

C-SHORE: HIGHER-ORDER VERIFICATION VIA COLLAPSIBLE PUSHDOWN SYSTEM SATURATION

CHRISTOPHER BROADBENT, ARNAUD CARAYOL, MATTHEW HAGUE, AND OLIVIER SERRE

Institut für Informatik (I7), Technische Universität München
e-mail address: broadben@in.tum.de

Laboratoire d'informatique de l'Institut Gaspard Monge, Université Paris-Est, and CNRS
e-mail address: Arnaud.Carayol@univ-mlv.fr

Department of Computer Science, Royal Holloway, University of London
e-mail address: matthew.hague@rhul.ac.uk

IRIF, Université Paris Diderot - Paris 7, and CNRS
e-mail address: olivier.serre@cnrs.fr

ABSTRACT. Higher-order recursion schemes (HORS) have received much attention as a useful abstraction of higher-order functional programs with a number of new verification techniques employing HORS model-checking as their centrepiece. We give an account of the C-SHORE tool, which contributed to the ongoing quest for a truly scalable model-checker for HORS by offering a different, automata theoretic perspective. C-SHORE implements the first practical model-checking algorithm that acts on a generalisation of pushdown automata equi-expressive with HORS called *collapsible pushdown systems* (CPDS). At its core is a backwards saturation algorithm for CPDS. Additionally, it is able to use information gathered from an approximate forward reachability analysis to guide its backward search. Moreover, it uses an algorithm that prunes the CPDS prior to model-checking and a method for extracting counter-examples in negative instances. We provide an up-to-date comparison of C-SHORE with the state-of-the-art verification tools for HORS. The tool and additional material are available from <http://cshore.cs.rhul.ac.uk>.

1. INTRODUCTION

Functional languages such as Haskell, OCaml and Scala strongly encourage the use of higher-order functions. This represents a challenge for software verification, which usually does not model recursion accurately, or models only first-order calls (e.g. SLAM [1] and Moped [28]). However, there has recently been much interest in a model called *higher-order recursion schemes (HORS)* (see e.g. [22]), which offers a way of abstracting functional programs in a manner that precisely models higher-order control-flow.

1998 ACM Subject Classification: F.1.1; Models of Computation; Automata.

Key words and phrases: Higher-Order; Verification; Model-Checking; Recursion Schemes; Collapsible Pushdown Systems; Saturation; Automata.

This article gives a full account of the C-SHORE tool [7] whose algorithms and implementation have been published in ICALP 2012 [3] and ICFP 2013 [4].

The execution trees of HORS enjoy decidable μ -calculus theories [22]. Even ‘reachability’ properties (subsumed by the μ -calculus) are very useful in practice. As a simple example, the safety of incomplete pattern matching clauses could be checked by asking whether the program can ‘reach a state’ where a pattern match failure occurs. More complex ‘reachability’ properties can be expressed using a finite automaton and could, for example, specify that the program respects a certain discipline when accessing a particular resource (see [15]). Despite even reachability being $(n - 1)$ -EXPTIME complete, recent research has revealed that useful properties of HORS can be checked in practice.

Kobayashi’s TRecS [19] tool, which checks properties expressible by a deterministic trivial Büchi automaton (all states accepting), was the first to achieve this. It works by determining whether a HORS is typable in an intersection-type system characterising the property to be checked [15]. In a bid to improve scalability, a number of other algorithms have subsequently been designed and implemented such as Kobayashi *et al.*’s GTRecS(2) [17, 20] and Neatherway *et al.*’s TravMC [21] tools, all based on intersection type inference. A recent overview of HORS model-checking was given by Ong [24].

This work is the basis of various techniques for verifying functional programs. In particular, Kobayashi *et al.* have developed MoChi [18] that checks safety properties of (OCaML) programs, and EHMTT Verifier [31] for tree processing programs. Both use a recursion schemes model-checker as a central component. Similarly, Ong and Ramsay [23] analyse programs with pattern matching employing recursion schemes as an abstraction.

Despite much progress, even the state-of-the-art TRecS does not scale to recursion schemes big enough to model realistically sized programs; achieving scalability while accurately tracking higher-order control-flow is a challenging problem. This article offers an automata-theoretic perspective on this challenge, providing a fresh set of tools that contrast with previous intersection-type approaches.

Collapsible pushdown systems (CPDS) [10] are an alternative representation of the class of execution trees that can be generated by recursion schemes (with linear-time mutual-translations between the two formalisms [10, 8]). While pushdown systems augment a finite-state machine with a stack and provide an ideal model for first-order programs [13], collapsible pushdown systems model higher-order programs by extending the stack of a pushdown system to a nested “stack-of-stacks” structure. The nested stack structure enables one to represent closures. Indeed the reader might find it helpful to view a CPDS as being a Krivine’s Abstract Machine in a guise making it amenable to the generalisation of techniques for pushdown model-checking. Salvati and Walukiewicz have studied in detail the connection with the Krivine abstract machine [27, 26].

For ordinary (‘order-1’) pushdown systems, a model-checking approach called *saturation* has been successfully implemented by tools such as Moped [28] and PDSolver [12]. Given a regular set of configurations of the pushdown system (represented by a finite automaton A acting on stacks), saturation can solve the ‘backward reachability problem’ by computing another finite automaton recognising a set of configurations from which a configuration in $\mathcal{L}(A)$ can be reached. This is a fixed-point computation that gradually adds transitions to A until it is ‘saturated’. If A recognises a set of error configurations, one can determine whether the pushdown system is ‘safe’ by checking if its initial configuration is recognised by the automaton computed by saturation.

The first contribution of this article was first presented in ICALP 2012. We extend the saturation method to a backward reachability analysis of collapsible pushdown systems [3].

This runs in PTIME when the number of control states is bounded. Crucially, this condition is satisfied when translating from recursion schemes of bounded arity with properties represented by automata of bounded size [10]. Whilst the HORS/intersection-type based tool GTRecS(2) also enjoys this fixed-parameter tractability, it times out on many benchmarks that our tool solves quickly. We remark also Ramsay *et al.* introduced a third fixed-parameter tractable algorithm in 2014 underlying their Preface tool [25].

Motivated by these facts, we then revisited the foundations of higher-order verification tools and introduced C-SHORE [7] — the first model-checking tool for the (direct) analysis of collapsible pushdown systems and that was presented in ICFP 2013 [4]. To achieve an efficient implementation, some substantial modifications and additions were made to the algorithm, leading to several novel practical and theoretical contributions:

- (1) An approximate *forward* reachability algorithm providing data
 - (a) ... allowing the CPDS to be pruned so that saturation receives a smaller input.
 - (b) ... employed by a modified saturation algorithm to guide its *backward* search.

This is essential for termination on most of our benchmarks.

- (2) A method for extracting witnesses to reachability.
- (3) A complete rework of the saturation algorithm to speed up fixed-point computation.
- (4) Experimental results comparing our approach with other tools.

We remark that the tools mentioned above propagate information forwards WRT the evaluation of the model. In contrast, the raw saturation algorithm works backwards, but we also show how forward and backward propagation can be combined.

Here we give a full account of the C-SHORE tool. This covers the saturation algorithm presented at ICALP 2012 as well as efficient algorithms implemented by C-SHORE in ICFP 2013. To prove soundness, we diverge from the ICALP 2012 proof, and instead base our proof on the witness generation algorithm presented in ICFP 2013. In particular, we present novel generalisations of witness generation, the forwards analysis, and the efficient fixed-point calculation to *alternating* CPDSs. These were only given for non-alternating CPDSs in ICFP 2013. The tool is available at <http://cshore.cs.rhul.ac.uk>.

Since C-SHORE was released, two new tools were released. Broadbent *et al.* introduced HorSat, which is an application of the saturation technique and initial forward analysis directly to intersection type analysis of HORS [6]. Recently HorSat2 improved the forwards analysis and made other algorithmic improvements [14]. Secondly, in POPL 2014, Ramsay *et al.* introduced Preface [25]. This is a type-based abstraction-refinement algorithm that attempts to simultaneously prove and disprove the property of interest. Both HorSat2 and Preface perform significantly better than previous tools.

Section 2 is an informal introduction to HORS and CPDS. In Section 3 we describe CPDS and how to represent sets of their configurations. The basic saturation algorithm introduced in ICALP 2012 is presented in Section 4 and proven correct in Section 5. Section 5.3 gives our generalised witness generation algorithm (that also implies soundness of saturation). We describe two optimisations to the saturation algorithm used by C-SHORE: an initial forwards analysis in Section 6 and an efficient fixed point computation in Section 7. Experimental results are in Section 8.

2. MODELLING HIGHER-ORDER PROGRAMS

In this section we give an informal introduction to the process of modelling higher-order programs for verification. In particular, we show how a simple example program can be




Figure 1: The behaviour of a toy recursion scheme.

modelled using a higher-order recursion scheme, and then we show how this scheme is evaluated using a collapsible pushdown system. For a more systematic approach to modelling higher-order programs with recursion schemes, we refer the reader to work by Kobayashi *et al.* [18]. This section is for background only, and can be safely skipped.

For this section, consider the toy example below.

```

Main = MakeReport Nil
MakeReport x = if * (Commit x)
               else (AddData x MakeReport)
AddData y f = if * (f Error) else (f Cons(_, y))

```

In this example, $*$ represents a non-deterministic choice (that may, for example, be a result of some input by the user). Execution begins at **Main** which aims to make a report which is a list. It sends an empty report to **MakeReport**. Either **MakeReport** finishes and commits the report somehow, or it adds an item to the head of the list using **AddData**, which takes the report so far, and a continuation. **AddData** either detects a problem with the new data (maybe it is inconsistent with the rest of the report) and flags an error by passing **Error** to the continuation, or extends the report with some item. In this case, since there is no error handling in **MakeReport**, an **Error** may be committed.

2.1. Higher-Order Recursion Schemes. We introduce, informally, higher-order recursion schemes. These are rewrite systems that generate the computation tree of a functional program. A rewrite rule takes the form

$$N \phi x \hookrightarrow t$$

where N is a (simply) typed non-terminal with (possibly higher-order) arguments ϕ and x . A term $N t_\phi t_x$ rewrites to t with t_ϕ substituted for ϕ and t_x substituted for x . Note that recursion schemes require t to be of ground type. We illustrate recursion schemes their use in analysis using the toy example from above. We can directly model our example with the scheme

$$\begin{aligned}
main &\hookrightarrow M \text{ nil} \\
M x &\hookrightarrow \text{or } (\text{commit } x) (A x M) \\
A y \phi &\hookrightarrow \text{or } (\phi \text{ error}) (\phi (\text{cons } y))
\end{aligned}$$

where M is the non-terminal associated with the **MakeReport** function, and A is the non-terminal associated with the **AddData** function; *nil*, *or*, *commit*, *error* and *cons* are terminal symbols of arity 0, 2, 1, 0 and 1 respectively (e.g. in the second rule, *or* takes the two arguments $(\text{commit } x)$ and $(A x M)$). The scheme above begins with the non-terminal *main* and, through a sequence of rewrite steps, generates a tree representation of the evolution of the program. Figure 1, described below, shows such a sequence.

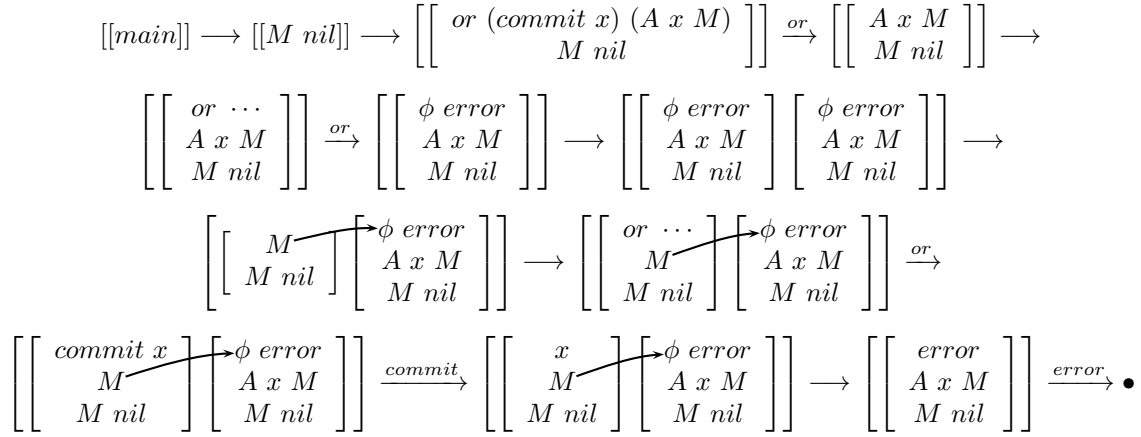


Figure 2: A stack evaluating the toy example.

Beginning with the non-terminal *main*, we apply the first rewrite rule to obtain the tree representing the term $(A \ nil)$. We then apply the second rewrite rule, instantiating x with *nil* to obtain the next tree in the sequence. This continues *ad infinitum* to produce a possibly infinite tree labelled only by terminals.

We aim to show the correctness of the program. I.e. the program never tries to *commit* an *error*. The rightmost tree in Figure 1, has a branch labelled *or, or, or, commit, error*. Note, *commit* is being called with an *error* report. In general we define the regular language $\mathcal{L}_{err} = or^* commit or^* error$. If the tree generated by the HORS contains a branch labelled by a word appearing in \mathcal{L}_{err} , then we have identified an error in the program.

2.2. Collapsible Pushdown Automata. Previous research into the verification of HORS has used *intersection types* (e.g. [16, 21]). Here we investigate a radically different approach exploiting the connection between HORS and an automata model called *collapsible pushdown automata* (CPDA). These two formalisms are, in fact, equivalent.

Theorem 2.1 (Equi-expressivity [10]). *For each order- n recursion scheme, there is an order- n collapsible pushdown automaton generating the same tree, and vice-versa. Furthermore, the translations in both directions are linear.* \square

We describe at a high level the structure of a CPDA and how they can be used to evaluate recursion schemes. In our case, this means outputting a sequence of non-terminals representing each path in the tree. More formal definitions are given in Section 3. At any moment, a CPDA is in a *configuration* $\langle p, w \rangle$, where p is a control state taken from a finite set \mathcal{P} , and w is a higher-order collapsible stack. In the following we will focus on the stack. Control states are only needed to ensure that sequences of stack operations occur in the correct order and are thus elided for clarity.

In our toy example, we have an order-2 HORS and hence an order-2 stack. An order-1 stack is a stack of characters from a finite alphabet Σ . An order-2 stack is a stack of order-1 stacks. Thus $[[main]]$ denotes the order-2 stack containing only the order-1 stack $[main]$; $[main]$ is an order-1 stack containing only the character *main*. In general Σ will contain all subterms appearing in the original statement of our toy example recursion scheme. The evolution of the CPDA stack is given in Figure 2 and explained below.

The first step is to rewrite *main* using $main \hookrightarrow M \text{ nil}$. Since $(M \text{ nil})$ is a subterm of our recursion scheme, we have $(M \text{ nil}) \in \Sigma$ and we rewrite the stack $[[main]]$ to $[[M \text{ nil}]]$. Next, we call M . As usual, a function call necessitates a new stack frame. In particular, we push the body of M (that is $(or \text{ (commit } x) (A \text{ } x \text{ } M)))$ onto the stack, giving the third stack in Figure 2. Note, we do not instantiate the variable x , hence we use only the subterms appearing in the recursion scheme.

Recall that we want to obtain a CPDA that outputs a sequence of terminals representing each path in the tree. To evaluate $or \text{ } (\dots) \text{ } (\dots)$ we output the terminal *or* and then (non-deterministically) choose a branch of the tree to follow. Let us choose $(A \text{ } x \text{ } M)$. Hence, the CPDA outputs *or* and rewrites the top term to $(A \text{ } x \text{ } M)$. Next we call A , pushing its body to the stack, then pick out the $(\phi \text{ error})$ branch of the *or* terminal. This takes us to the beginning of the second row of Figure 2.

To proceed, we evaluate $(\phi \text{ error})$. To do this, we have to know the value of ϕ . We can obtain this information by inspecting the stack and seeing that the second argument of the call of A is M . However, since we can only see the top of a stack, we would have to remove the character $(\phi \text{ error})$ to determine that $\phi = M$, thus losing our place in the computation.

However, an order-2 stack is able — via a $push_2$ operation — to create a copy of its topmost order-1 stack. After this copy (note that the top of the stack is written on the left) we delve into the copy of the stack to find the value of ϕ . Simultaneously we create a *collapse link*, pictured as an arrow from M to the term $(\phi \text{ error})$. This collapse link points from M to the context in which M will be evaluated. In particular, if we need to know the value of x in the body of M , we need to know that M was called with the *error* argument, within the term $(\phi \text{ error})$; the collapse link points to this information (i.e. encodes a closure in the stack). We can access this information via a *collapse* operation. These are the two main features of a higher-order collapsible stack, described formally in the next section.

To continue, we push the body of M on to the stack, output the *or* symbol and choose the $(\text{commit } x)$ branch. Since *commit* is a terminal, we output it and evaluate x . To compute x , we look into the stack and follow the collapse link from M to $(\phi \text{ error})$. We do not create a copy of the stack here because x is an order-0 variable and thus represents a self-contained execution. Since x has value *error*, we output it and terminate. This completes the execution corresponding to the error branch identified in Figure 1.

2.3. Collapsible Pushdown Systems. The CPDA output *or, or, or, commit, error* in the execution above. This is an error sequence in \mathcal{L}_{err} and should be flagged. In general, we take the finite automaton A representing the regular language \mathcal{L}_{err} and form a synchronised product with the CPDA. This results in a CPDA that does not output any symbols, but instead keeps in its control state the progression of A . Thus we are interested in whether the CPDA is able to reach an accepting state of A , not the language it generates. We call a CPDA without output symbols a *collapsible pushdown system* (CPDS), and the question of whether a CPDS can reach a given state is the reachability problem. This is the subject of the remainder of the paper.


3. PRELIMINARIES

3.1. Collapsible Pushdown Systems. We give the definition of higher-order collapsible stacks and their operations, before giving the definition of collapsible pushdown systems.

3.1.1. Higher-Order Collapsible Stacks. Higher-order collapsible stacks are a nested “stack-of-stacks” structure over a stack alphabet Σ . Each stack character contains a pointer — called a “link” — to a position lower down in the stack. The stack operations defined below, create copies of sub-stacks. The link is intuitively a pointer to the context in which the stack character was first created. We define collapse links formally below.

Definition 3.1 (Order- n Collapsible Stacks). Given a finite set of stack characters Σ , an *order-0 stack* is simply a character $a \in \Sigma$. An *order- n stack* is a sequence $w = [w_1 \dots w_\ell]_n$ such that each w_i is an order- $(n-1)$ stack and each character on the stack is augmented with a collapse link. Let $Stacks_n$ denote the set of order- n stacks.

An order- n stack can be represented naturally as an edge-labelled word-graph over the alphabet $\{[_{n-1}, \dots, [_{1, 1}, \dots,]_{n-1}] \uplus \Sigma$, with additional collapse-links pointing from a stack character in Σ to the beginning of the graph representing the target of the link. An example order-3 stack is given below, with only a few collapse links shown. The collapse links range from order-3 to order-1 respectively.



Given an order- n stack $[w_1 \dots w_\ell]_n$, we define

$$\begin{aligned} top_n([w_1 \dots w_\ell]_n) &= w_1 && \text{when } \ell > 0 \\ top_n([]_n) &= []_{n-1} && \text{otherwise} \\ top_k([w_1 \dots w_\ell]_n) &= top_k(w_1) && \text{when } k < n \text{ and } \ell > 0 \end{aligned}$$

noting that $top_k(w)$ is undefined if $top_{k'}(w)$ is empty for any $k' > k$. We also remove the top portion of a top_k stack using

$$\begin{aligned} bot_n^i([w_1 \dots w_\ell]_n) &= [w_{\ell-i+1} \dots w_\ell]_n && \text{when } i \leq \ell \text{ and } \ell > 0 \\ bot_k^i([w_1 \dots w_\ell]_n) &= [bot_k^i(w_1)w_2 \dots w_\ell]_n && \text{when } k < n \text{ and } \ell > 0. \end{aligned}$$

Formally, a collapse link is a pair (k, i) where $1 \leq k \leq n$ and $i > 0$. For $top_1(w) = a$ where a has the link (k, i) , the destination of the link is $bot_k^i(w)$. We disallow collapse links where bot_k^i is undefined. The example stack above is $[[[a^{(3,1)}b]_1]_2[[c^{(2,1)}]_1]_2[[d^{(1,1)}e]_1]_2]_3$, where superscripts denote collapse links. Often, we omit these superscripts for readability.

3.1.2. Operations on Order- n Collapsible Stacks. The following operations may be performed on an order- n collapsible stack.

$$\begin{aligned} \mathcal{O}_n &= \{pop_1, \dots, pop_n\} \cup \{push_2, \dots, push_n\} \cup \\ &\quad \{collapse_2, \dots, collapse_n\} \cup \{push_a^1, \dots, push_a^n, rew_a \mid a \in \Sigma\} \end{aligned}$$

We say $o \in \mathcal{O}_n$ is of order- k when k is minimal such that $o \in \mathcal{O}_k$. E.g., $push_k$ is of order k .

The $collapse_k$ operation is non-standard in the sense of Hague *et al.* [10] and has the semantics of a normal collapse, with the additional constraint that the top character has an order- k link. The standard version of collapse can be simulated with a non-deterministic choice on the order of the stack link. In the other direction, we can store in the stack alphabet the order of the collapse link attached to each character on the stack.

When u is a k -stack and $v = [v_1 \dots v_\ell]_n$ is an n -stack with $k \leq n$, we define $u :_k v$ as the stack obtained by adding u on top of the topmost $(k+1)$ -stack of v . Formally, we let

$$\begin{aligned} u :_k v &= [uv_1 \dots v_\ell]_n && \text{when } k = n \\ u :_k v &= [(u :_k v_1)v_2 \dots v_\ell]_n && \text{when } k < n \end{aligned}$$

We define each stack operation in turn for an order- n stack w . Collapse links are created by the $push_a^k$ operations, which add a character to the top of a given stack w with a link pointing to $pop_k(w)$.

- (1) We set $pop_k(w) = v$ when w decomposes into $u :_k v$ for a non-empty u .
- (2) We set $push_k(w) = u :_k u :_k v$ when $w = u :_k v$.
- (3) We set $collapse_k(w) = bot_k^i(w)$ where $top_1(w) = a^{(k,i)}$ for some i .
- (4) We set $push_b^k(w) = b^{(k,\ell-1)} :_1 w$ where $top_{k+1}(w) = [w_1 \dots w_\ell]_{k+1}$.
- (5) We set $rew_b(w) = b^{(k,i)} :_1 v$ where $w = a^{(k,i)} :_1 v$.

Note that, for a $push_k$ operation, links outside of $u = top_k(w)$ point to the same destination in both copies of u , while links pointing within u point within the respective copies of u . For full introduction, we refer the reader to Hague *et al.* [10]. In Section 4.3 we give several example stacks and show how the stack operations affect them.

3.1.3. Collapsible Pushdown Systems. We define alternating collapsible pushdown systems.

Definition 3.2 (Collapsible Pushdown Systems). An alternating order- n *collapsible pushdown system* (*collapsible PDS*) is a tuple $\mathcal{C} = (\mathcal{P}, \Sigma, \mathcal{R})$ where \mathcal{P} is a finite set of control states, Σ is a finite stack alphabet, and $\mathcal{R} \subseteq (\mathcal{P} \times \Sigma \times \mathcal{O}_n \times \mathcal{P}) \cup (\mathcal{P} \times 2^{\mathcal{P}})$ is a set of rules.

We write *configurations* of a collapsible PDS as a pair $\langle p, w \rangle$ where $p \in \mathcal{P}$ and $w \in Stacks_n$. We write $\langle p, w \rangle \longrightarrow \langle p', w' \rangle$ to denote a transition from a rule (p, a, o, p') with $top_1(w) = a$ and $w' = o(w)$. Furthermore, we have a transition $\langle p, w \rangle \longrightarrow \{ \langle p', w \rangle \mid p' \in P \}$ whenever we have a rule $p \rightarrow P$. A non-alternating collapsible PDS has no rules of this second form. We write C to denote a set of configurations.

3.2. Regularity of Collapsible Stacks. We will present an algorithm that operates on sets of configurations. For this we use order- n stack automata, thus defining a notion of regular sets of stacks. These have a nested structure based on a similar automata model by Bouajjani and Meyer [2]. The handling of collapse links is similar to automata introduced by Broadbent *et al.* [5], except we read stacks top-down rather than bottom-up.

Definition 3.3 (Order- n Stack Automata). An *order- n stack automaton*

$$A = (\mathbb{Q}_n, \dots, \mathbb{Q}_1, \Sigma, \Delta_n, \dots, \Delta_1, \mathcal{F}_n, \dots, \mathcal{F}_1)$$


is a tuple where Σ is a finite stack alphabet, and

- (1) for all $n \geq k \geq 2$, we have \mathbb{Q}_k is a finite set of states, $\Delta_k \subseteq \mathbb{Q}_k \times \mathbb{Q}_{k-1} \times 2^{\mathbb{Q}_k}$ is a transition relation, and $\mathcal{F}_k \subseteq \mathbb{Q}_k$ is a set of accepting states, and
- (2) \mathbb{Q}_1 is a finite set of states, $\Delta_1 \subseteq \bigcup_{2 \leq k \leq n} (\mathbb{Q}_1 \times \Sigma \times 2^{\mathbb{Q}_k} \times 2^{\mathbb{Q}_1})$ a transition relation, and $\mathcal{F}_1 \subseteq \mathbb{Q}_1$ a set of accepting states.

Stack automata are alternating automata that read the stack in a nested fashion. Order- k stacks are recognised from states in \mathbb{Q}_k . A transition $(q, q', Q) \in \Delta_k$ from q to Q for some $k > 1$ can be fired when the top_{k-1} stack is accepted from $q' \in \mathbb{Q}_{(k-1)}$. The remainder of the stack must be accepted from all states in Q . At order-1, a transition (q, a, Q_{col}, Q) is a standard alternating a -transition with the additional requirement that the stack pointed to by the collapse link of a is accepted from all states in Q_{col} . A stack is accepted if a subset of \mathcal{F}_k is reached at the end of each order- k stack. In the next section, we formally define the runs of a stack automaton. We write $w \in \mathcal{L}_q(A)$ whenever w is accepted from a state q . For

ease of presentation, we write $q \xrightarrow{q'} Q \in \Delta_k$ instead of $(q, q', Q) \in \Delta_k$ and $q \xrightarrow[a]{Q_{col}} Q \in \Delta_1$ instead of $(q, a, Q_{col}, Q) \in \Delta_1$.

A (partial) run is pictured below, using $q_3 \xrightarrow{q_2} Q_3 \in \Delta_3, q_2 \xrightarrow{q_1} Q_2 \in \Delta_2$ and $q_1 \xrightarrow[a]{Q_{col}} Q_1 \in \Delta_1$. The node labelled Q_{col} begins a run on the stack pointed to by the collapse link of a . Note that the label of this node may contain other elements apart from Q_{col} . These additional elements come from the part of the run coming from the previous node (and other collapse links).



3.2.1. Formal Definition of a Run. For $n \geq k > 1$, we write $Q_1 \xrightarrow{Q'} Q_2$ to denote an order- k transition from a set of states whenever $Q_1 = \{q_1, \dots, q_\ell\}$ and for each $1 \leq i \leq \ell$ we have $q_i \xrightarrow{q'_i} Q_i$ and $Q' = \{q'_1, \dots, q'_\ell\}$ and $Q_2 = \bigcup_{1 \leq i \leq \ell} Q_i$. The analogous notation at order-1 is a special case of the short-form notation defined in Section 4.1.

Fix a stack automaton

$$A = (\mathbb{Q}_n, \dots, \mathbb{Q}_1, \Sigma, \Delta_n, \dots, \Delta_1, \mathcal{F}_n, \dots, \mathcal{F}_1) .$$

Recall the word-graph representation of a stack. We say a node *contains* a character if its exiting edge is labelled by the character.

A stack (word-graph) w is accepted by A from states $Q_0 \subseteq \mathbb{Q}_k$ — written $w \in \mathcal{L}_{Q_0}(A)$ — whenever the nodes of the graph can be labelled by elements of $\bigcup_{1 \leq k' \leq n} 2^{\mathbb{Q}_{k'}}$ such that

- (1) Q_0 is a subset of the label of the node containing the first $[_{k-1}$ character of the word, or if $k = 1$, the first character $a \in \Sigma$, and
- (2) for any node containing a character $[_{k'}$ labelled by Q , then for all $q_1 \in Q$, there exists some transition $(q_1, q_2, Q_1) \in \Delta_{k'+1}$ such that q_2 appears in the label of the succeeding node and Q_1 is a subset of the label of the node succeeding the matching $]_{k'}$ character, and
- (3) for any node containing a character $]_{k'}$, the label Q is a subset of $\mathcal{F}_{k'}$, and the final node of an order- k stack is labelled by $Q \subseteq \mathcal{F}_k$, and
- (4) for any node containing a character $a \in \Sigma$, labelled by Q , for all $q' \in Q$, there exists some transition $(q', a, Q_{col}, Q') \in \Delta_1$ such that Q_{col} is a subset of the label of the node reachable from the collapse branch of the edge labelled a , and Q' is a subset of the label of the succeeding node.

That is, a stack automaton is essentially a stack- and collapse-aware alternating automaton, where collapse links are treated as special cases of the alternation.

3.2.2. Properties of Stack Automata. We show here that stack automata membership is linear time, while emptiness is PSPACE. These are standard complexities for alternating finite word automata.

Several further results can also be shown [3], which we have omitted for space reasons: the sets of stacks accepted by these automata form an effective Boolean algebra (note that complementation causes a blow-up in the size of the automaton); and they accept the same

family of collapsible stacks as the automata used by Broadbent *et al.* [5]. We omit these here for space reasons.

Proposition 3.4 (Stack Automata Membership). *Membership of order- n stack automata can be tested in linear time in the size of the input stack and stack automaton.*

Proof. Take a stack w and let

$$A = (\mathbb{Q}_n, \dots, \mathbb{Q}_1, \Sigma, \Delta_n, \dots, \Delta_1, \mathcal{F}_n, \dots, \mathcal{F}_1) .$$

The membership algorithm iterates from the bottom (end) of the stack to the top (beginning). Take the graph representing w . We start by labelling the final node with the set \mathcal{F}_n . It is easy to verify at all stages that if the label of a node contains a state q , then the stack from that node is in $\mathcal{L}_q(A)$. Now, suppose we have labelled up to a given node. For convenience, we will refer to nodes by their labelled state set, and we show, by cases, how to label the preceding node with a state-set Q_0 .

In the first case, we have node Q_1 connected to Q_0 by a $]_k$ character. That is

$$Q_0 \xrightarrow{]}_k Q_1 \quad \dots$$

where Q_0 is the set \mathcal{F}_k . This is because Q_0 labels the end of an order- k stack.

In the next case the connection is by a character a with a collapse link and we have

$$Q_0 \xrightarrow{a} Q_1 \quad \dots \quad Q_{col} \quad \dots$$

where $Q_0 = \left\{ q \mid q \xrightarrow[Q'_{col}]{a} Q'_1 \in \Delta_1 \wedge Q'_1 \subseteq Q_1 \wedge Q'_{col} \subseteq Q_{col} \right\}$. Thus, any state in Q_0 has a transition to states from which the remainder of the stack is accepted.

In the final case we have

$$Q_0 \xrightarrow{[}_k Q_1 \quad \dots \xrightarrow{]}_k Q_2 \quad \dots$$

where $]_k$ matches $[_k$. We define $Q_0 = \left\{ q \mid q \xrightarrow{q'} Q \in \Delta_{k+1} \wedge q' \in Q_1 \wedge Q \subseteq Q_2 \right\}$. Thus,

there is a state $q \in Q_0$ whenever there is a transition $q \xrightarrow{q'} Q$ such that the next order- k stack is accepted from q' and the remainder of the stack is accepted from Q .

Thus, after this labelling, we can test whether $w \in \mathcal{L}_{Q_0}(A)$ for $Q_0 \subseteq \mathbb{Q}_k$ by checking whether $Q_0 \subseteq Q$ where Q labels the node containing the first $]_{k-1}$ character of the word, or if $k = 1$, the first character $a \in \Sigma$. \square

We now show that emptiness checking is PSPACE-complete.

Proposition 3.5 (Emptiness of Stack Automata). *Given an order- n stack automaton A and a state q of A , deciding if there is some $w \in \mathcal{L}_q(A)$ is PSPACE-complete.*

Proof. The problem is already PSPACE-hard for alternating word automata which correspond to the order-1 case [9]. We now establish the upperbound.

Let $A = (\mathbb{Q}_n, \dots, \mathbb{Q}_1, \Sigma, \Delta_n, \dots, \Delta_1, \mathcal{F}_n, \dots, \mathcal{F}_1)$ be a stack automaton. We use a fixed point to compute the set $\mathbb{Q}_{ac} := \{ q \in \mathbb{Q}_k \mid \mathcal{L}_q(A) \neq \emptyset \}$.

For this, we take Q_0 to be the set of all final states. For all $i \geq 0$, we define Q_{i+1} by adding to Q_i all states $q \in \mathbb{Q}_k$ such that:

- for $k > 1$, there exists a transition $q \xrightarrow{q'} Q \in \Delta_k$ with $Q \subseteq Q_i$ and $q' \in Q_i$.
- for $k = 1$, there exists a transition $q \xrightarrow[Q_{col}]{a} Q \in \Delta_1$ with $Q, Q_{col} \subseteq Q_i$.

As the sequence of the Q_i is increasing (for inclusion), we have $Q_{j+1} = Q_j$ for some index $j \geq 0$. We claim that Q_j is the set Q_{ac} . A straightforward induction shows that for all i we have $Q_i \subseteq Q_{ac}$ and hence $Q_j \subseteq Q_{ac}$.

Assume toward a contradiction, that the converse inclusion does not hold. Let w be smallest stack accepted by a state $q \in \mathbb{Q}_k \setminus Q_j$. If $k = 1$, then $w = a^{(k,i')} :_1 w'$ and there exists a transition $q \xrightarrow[Q_{col}]{a} Q \in \Delta_1$ with w' accepted from Q and u is accepted from Q_{col} where $u = \text{collapse}_k(w)$. As q does not belong to Q_j , then either u or w' is accepted from a state not in Q_j . As u and w' are both smaller than w , this contradicts the definition of w . The case $k > 1$ is similar.

The algorithm to test emptiness from a given state $p \in \mathbb{Q}_n$ consists of computing Q_j by iteratively computing the Q_i and then checking if p belongs to Q_j . \square

4. SATURATION ALGORITHM

Given a CPDS \mathcal{C} and a stack automaton A_0 with a state $q_p \in \mathbb{Q}_n$ for each control state p in \mathcal{C} , let $Pre_{\mathcal{C}}^*(A_0) = \bigcup_{\alpha < \omega} Pre_{\mathcal{C}}^\alpha(A_0)$ where

$$\begin{aligned} Pre_{\mathcal{C}}^0(A_0) &= \{ \langle p, w \rangle \mid w \in \mathcal{L}_{q_p}(A_0) \} \\ Pre_{\mathcal{C}}^{\alpha+1}(A_0) &= \left\{ \langle p, w \rangle \mid \begin{array}{l} \exists \langle p', w' \rangle \in Pre_{\mathcal{C}}^\alpha(A_0) \vee \\ \exists \langle p, w \rangle \longrightarrow C \subseteq Pre_{\mathcal{C}}^\alpha(A_0) \end{array} \right\} \end{aligned}$$

recalling that C denotes a set of configurations. We build a stack automaton recognising $Pre_{\mathcal{C}}^*(A_0)$. We begin with A_0 and iterate a saturation function denoted Π — which adds new transitions to A_0 — until a ‘fixed point’ has been reached. That is, we iterate $A_{i+1} = \Pi(A_i)$ until $A_{i+1} = A_i$. As the number of states is bounded, we eventually obtain this, giving us the following theorem.

Theorem 4.1. *Given an alternating CPDS \mathcal{C} and a stack automaton A_0 , we can construct an automaton A accepting $Pre_{\mathcal{C}}^*(A_0)$.* \square

The construction runs in n -EXPTIME for alternating CPDS — which is optimal — and can be improved to $(n-1)$ -EXPTIME for non-alternating CPDS when the initial automaton satisfies a certain notion of *non-alternation*, again optimal. Correctness and complexity are discussed in subsequent sections.

4.1. Notation and Conventions.

Number of Transitions. We assume for all $q \in \mathbb{Q}_k$ and $Q \subseteq \mathbb{Q}_k$ that there is at most one transition of the form $q \xrightarrow{q'} Q \in \Delta_k$. This condition can easily be ensured on A_0 by replacing pairs of transitions $q \xrightarrow{q_1} Q$ and $q \xrightarrow{q_2} Q$ with a single transition $q \xrightarrow{q'} Q$, where q' accepts the union of the languages of stacks accepted from q_1 and q_2 . The construction maintains this condition.

Short-form Notation. We introduce some short-form notation for runs. Consider the example run in Section 3.2. In this case, we write $q_3 \xrightarrow[Q_{col}]{a} (Q_1, Q_2, Q_3)$, $q_3 \xrightarrow{q_1} (Q_2, Q_3)$, and $q_3 \xrightarrow{q_2} (Q_3)$. In general, we write

$$q \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_k) \text{ and } q \xrightarrow{q'} (Q_{k'+1}, \dots, Q_k).$$

In the first case, $q \in \mathbb{Q}_k$ and there exist q_{k-1}, \dots, q_1 such that $q \xrightarrow{q_{k-1}} Q_k \in \Delta_k$, $q_{k-1} \xrightarrow{q_{k-2}} Q_{k-1} \in \Delta_{k-1}, \dots, q_1 \xrightarrow[Q_{col}]{a} Q_1 \in \Delta_1$. Thus, we capture nested sequences of initial transitions from q . Since we assume at most one transition between any state and set of states, the intermediate states q_{k-1}, \dots, q_1 are uniquely determined by q, a, Q_{col} and Q_1, \dots, Q_k .

In the second case $q \in \mathbb{Q}_k$, $q' \in \mathbb{Q}_{k'}$, and there exist $q_{k-1}, \dots, q_{k'+1}$ with $q \xrightarrow{q_{k-1}} Q_k \in \Delta_k$, $q_{k-1} \xrightarrow{q_{k-2}} Q_{k-1} \in \Delta_{k-1}, \dots, q_{k'+2} \xrightarrow{q_{k'+1}} Q_{k'+2} \in \Delta_{k'+2}$ and $q_{k'+1} \xrightarrow{q'} Q_{k'+1} \in \Delta_{k'+1}$.

We lift the short-form transition notation to transitions from sets of states. We safely assume that state-sets $\mathbb{Q}_n, \dots, \mathbb{Q}_1$ are disjoint. Suppose $Q = \{q_1, \dots, q_\ell\}$ and for all $1 \leq i \leq \ell$ we have $q_i \xrightarrow[Q_{col}^i]{a} (Q_1^i, \dots, Q_k^i)$. Then we have $Q \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_k)$ where $Q_{col} = \bigcup_{1 \leq i \leq \ell} Q_{col}^i$ and for all k , $Q_k = \bigcup_{1 \leq i \leq \ell} Q_k^i$. Because a collapse link can only be of one order, we insist that $Q_{col} \subseteq \mathbb{Q}_k$ for some k .

Note that a transition to the empty set is distinct from having no transition.

Initial States. We say a state is *initial* if it is of the form $q_p \in Q_n$ for some control state p or if it is a state $q_k \in Q_k$ for $k < n$ such that there exists a transition $q_{k+1} \xrightarrow{q_k} Q_{k+1}$ in Δ_{k+1} . We make the assumption that all initial states do not have any incoming transitions and that they are not final¹.

Adding Transitions. Finally, when we add a transition $q_n \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ to the automaton, then for each $n \geq k > 1$, we add $q_k \xrightarrow{q_{k-1}} Q_k$ to Δ_k (if a transition between q_k and Q_k does not already exist, otherwise we use the existing transition and state q_{k-1}) and add $q_1 \xrightarrow[Q_{col}]{a} Q_1$ to Δ_1 .

Justified Transitions. When we add transitions via the saturation function we also add *justifications* to the new transitions that are not derived from alternating transitions. These justifications indicate the provenance of each new transition. This later permits counter example generation for CPDSs, as shown in Section 5.3.

To each $t = q \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ we will define the justification $J(t)$ to be either 0 (indicating the transition is in A_0), a pair (r, i) , a tuple (r, t', i) , (r, T, i) or a tuple (r, t', T, i) where r is a rule of the CPDS, i is the number of iterations saturation required to introduce the transition, t' is a long-form transition and T is a set of such transitions.

Note that we apply J to the short-form notation. In reality, we associate each justification with the unique order-1 transition $q_1 \xrightarrow[Q_{col}]{a} Q_1$ associated to each t .

¹Hence automata cannot accept empty stacks from initial states. This can be overcome by introducing a bottom-of-stack symbol.

4.2. The Saturation Function. We are now ready to give the saturation function Π for a given $\mathcal{C} = (\mathcal{P}, \Sigma, \mathcal{R})$. As described above, we apply this function to A_0 until a fixed point is reached. First set $J(t) = 0$ for all transitions of A_0 . The intuition behind the saturation rules can be quickly understood via a rewrite rule (p, a, rew_b, p') which leads to the addition of a transition $q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ whenever there already existed a transition $q_{p'} \xrightarrow[Q_{col}]{b} (Q_1, \dots, Q_n)$. Because the rewrite can change the control state from p to p' and the top character from a to b , we must have an accepting run from q_p with a on top whenever we had an accepting run from $q_{p'}$ with b on top. We give examples and intuition of the more complex steps in Section 4.3, which may be read alongside the definition below.

Definition 4.2 (The Saturation Function Π). Given an order- n stack automaton A_i we define $A_{i+1} = \Pi(A_i)$. The state-sets of A_{i+1} are defined implicitly by the transitions which are those in A_i plus, for each $r = (p, a, o, p') \in \mathcal{R}$,

- (1) when $o = pop_k$, for each $q_{p'} \xrightarrow{q_k} (Q_{k+1}, \dots, Q_n)$ in A_i , add

$$t = q_p \xrightarrow[\emptyset]{a} (\emptyset, \dots, \emptyset, \{q_k\}, Q_{k+1}, \dots, Q_n)$$

to A_{i+1} and set $J(t) = (r, i+1)$ whenever t is not already in A_{i+1} ,

- (2) when $o = push_k$, for each $t = q_{p'} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_k, \dots, Q_n)$ and T of the form $Q_k \xrightarrow[Q'_{col}]{a} (Q'_1, \dots, Q'_k)$ in A_i , add to A_{i+1} the transition

$$t' = q_p \xrightarrow[Q_{col} \cup Q'_{col}]{a} \begin{pmatrix} Q_1 \cup Q'_1, \dots, Q_{k-1} \cup Q'_{k-1}, \\ Q'_k, \\ Q_{k+1}, \dots, Q_n \end{pmatrix}$$

and set $J(t') = (r, t, T, i+1)$ if t' is not already in A_{i+1} ,

- (3) when $o = collapse_k$, when $k = n$, add $t = q_p \xrightarrow[\{q_{p'}\}]{a} (\emptyset, \dots, \emptyset)$ if it does not exist, and

when $k < n$, for each transition $q_{p'} \xrightarrow{q_k} (Q_{k+1}, \dots, Q_n)$ in A_i , add to A_{i+1} the transition $t = q_p \xrightarrow[\{q_k\}]{a} (\emptyset, \dots, \emptyset, Q_{k+1}, \dots, Q_n)$ if t does not already exist. In all cases, if t is added, set $J(t) = (r, i+1)$,

- (4) when $o = push_b^k$ for all transitions $t = q_{p'} \xrightarrow[Q_{col}]{b} (Q_1, \dots, Q_n)$ and $T = Q_1 \xrightarrow[Q'_{col}]{a} Q'_1$ in A_i with $Q_{col} \subseteq Q_k$, add to A_{i+1} the transition

$$t' = q_p \xrightarrow[Q'_{col}]{a} (Q'_1, Q_2, \dots, Q_k \cup Q_{col}, \dots, Q_n) ,$$

and set $J(t') = (r, t, T, i+1)$ if t' is not already in A_{i+1} ,

- (5) when $o = rew_b$ for each transition $t = q_{p'} \xrightarrow[Q_{col}]{b} (Q_1, \dots, Q_n)$ in A_i , add to A_{i+1} the transition $t' = q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$, setting $J(t') = (r, t, i)$ when t' is not already in A_{i+1} .

Finally, for every rule $p \rightarrow P$, let $Q = \{q_{p'} \mid p' \in P\}$, then, for each T of the form $Q \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$, add a transition $q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ and set $J(t') = (r, T, i+1)$ if t' is not already in A_{i+1} .




Figure 3: A sequence of saturation steps.

From A_0 , we iterate $A_{i+1} = \Pi(A_i)$ until $A_{i+1} = A_i$. Generally, as we show in Proposition 5.2, we terminate in n -EXPTIME. When A_0 satisfies a “non-alternating” property (e.g. when we are only interested in reaching a designated control state), we can restrict Π to only add transitions where Q_n has at most one element, giving $(n - 1)$ -EXPTIME complexity. In all cases saturation is linear in the size of Σ .

4.3. Examples of Saturation. As an example, consider a CPDS with the run

$$\langle p_1, [b] [c] [d] \rangle \xrightarrow{push_a^2} \langle p_2, [\overline{ab}] [c] [d] \rangle \xrightarrow{push_2} \langle p_3, [\overline{ab}] [\overline{ab}] [c] [d] \rangle \xrightarrow{collapse_2} \langle p_4, [c] [d] \rangle \xrightarrow{pop_2} \langle p_5, [d] \rangle.$$

Figure 3 shows the sequence of saturation steps, beginning with an accepting run of $\langle p_5, [d] \rangle$ and finishing with an accepting run of $\langle p_1, [b] [c] [d] \rangle$. The individual steps are explained below.

Initial Automaton. The top of Figure 3 shows a stack automaton containing the transitions $q_{p5} \xrightarrow{q_1} \emptyset$ and $q_1 \xrightarrow{d} \emptyset$, which we write $q_{p5} \xrightarrow{d} (\emptyset, \emptyset)$. This gives the run over $\langle p_5, [d] \rangle$.

Rule (p_4, c, pop_2, p_5) . When the saturation step considers such a pop rule, it adds $q_{p4} \xrightarrow{c} (\emptyset, \{q_{p5}\})$. We add such a transition because we only require the top order-1 stack (removed by pop_2) to have the top character c (hence \emptyset is the next order-1 label), and after the pop_2 the remaining stack needs to be accepted from q_{p5} (hence $\{q_{p5}\}$ is the next order-2 label). This new transition allows us to construct the next run over $\langle p_4, [c] [d] \rangle$ in Figure 3.

Rule $(p_3, a, collapse_2, p_4)$. Similarly to the pop rule above, the saturation step adds the transition $q_{p3} \xrightarrow{a} (\emptyset, \emptyset)$. The addition of such a transition allows us to construct the

pictured run over $\langle p_3, [ab] [ab] [c] [d] \rangle$ (collapse links omitted), recalling that $\emptyset \xrightarrow{\emptyset} \emptyset$, $\emptyset \xrightarrow{a} \emptyset$ and $\emptyset \xrightarrow{b} \emptyset$ transitions are always possible due to the empty initial set. Note that the labelling of $\{q_{p4}\}$ comes from the collapse link on the topmost a character on the stack.

Rule $(p_2, a, push_2, p_3)$. Consider the run from q_{p_3} in Figure 3. The initial transition of the run accepting the first order-1 stack is $q_{p_3} \xrightarrow[\{q_{p_4}\}]{a} (\emptyset, \emptyset)$. We also have $\emptyset \xrightarrow{\emptyset} \emptyset$ (trivially) accepting the second order-1 stack. Any $push_2$ predecessor of this stack must have a top order-1 stack that could have appeared twice at the top of the stack from q_{p_3} . Thus, the saturation step makes the intersection of the initial order-1 transitions of first two order-1 stacks. This results in the transition $q_{p_2} \xrightarrow[\{q_{p_4}\} \cup \emptyset]{a} (\emptyset \cup \emptyset, \emptyset)$, which is used to form the shown run over $\langle p_2, [ab] [c] [d] \rangle$ (collapse links omitted).

Rule $(p_1, b, push_a^2, p_2)$. The run from q_{p_2} in Figure 3 begins with $q_{p_2} \xrightarrow[\{q_{p_4}\}]{a} (\emptyset, \emptyset)$ and $\emptyset \xrightarrow[\emptyset]{b} \emptyset$. Note that the $push_a^2$ gives a stack with ab on top. Moreover, the collapse link on a should point to the order-1 stack just below the current top one. Since the transition from q_{p_2} requires that the linked-to stack is accepted from q_{p_4} , we need this requirement in the preceding stack (accepted from q_{p_1} and without the a on top). Thus, we move the target of the collapse link into the order-2 destination of the new transition. That is, for $push_a^2$ we create $q_{p_1} \xrightarrow[\emptyset]{b} (\emptyset, \emptyset \cup \{q_{p_4}\})$. From this we can construct an accepting run over $\langle p_1, [b] [c] [d] \rangle$.

5. CORRECTNESS AND COMPLEXITY

In this section we show the complexity and correctness of saturation. We prove soundness by a witness generation algorithm. This is an extension of the witness generation given in ICFP 2013 [4] to the case of alternating CPDSs. In ICALP 2012 [3] we gave a more denotational proof of soundness which worked by showing that all transitions added by saturation respect the “meaning” of the transitions in the automata representing $Pre_{\mathcal{C}}^*(A_0)$. This proof used a slightly different formulation of CPDS as Annotated Pushdown Systems. Although we believe this soundness proof to be more elegant, we do not repeat it here for space reasons (since it would require the definition of annotated pushdown systems).

Theorem 5.1. *For a given \mathcal{C} and A_0 , let $A = A_i$ where i is the least index such that $A_{i+1} = \Pi(A_i)$. We have $w \in \mathcal{L}_{q_p}(A)$ iff $\langle p, w \rangle \in Pre_{\mathcal{C}}^*(A_0)$. \square*

The proof is given in the following sections. Completeness is by a straightforward induction over the “distance” to A_0 . Soundness is the key technical challenge.

Proposition 5.2. *The saturation construction for an alternating order- n collapsible PDS \mathcal{C} and an order- n stack automaton A_0 runs in n -EXPTIME, which is optimal.*

Proof. Let $2 \uparrow_0 (\ell) = \ell$ and $2 \uparrow_{i+1} (\ell) = 2^{2 \uparrow_i (\ell)}$. The number of states of A is bounded by $2 \uparrow_{(n-1)} (\ell)$ where ℓ is the size of \mathcal{C} and A_0 : each state in \mathbb{Q}_k was either in A_0 or comes from a transition in Δ_{k+1} . Since the automata are alternating, there is an exponential blow up at each order except at order- n . Each iteration of the algorithm adds at least one new transition. Only $2 \uparrow_n (\ell)$ transitions can be added. Since reachability for alternating higher-order pushdown systems is complete for n -EXPTIME [11], our algorithm is optimal. \square


The complexity of reachability for non-alternating collapsible PDS is in $(n-1)$ -EXPTIME. The cause of the additional exponential blow up is in the alternation of the stack automata. However, for a suitable notion of *non-alternating* stack automata, our algorithm can be adapted to run in $(n-1)$ -EXPTIME, when the CPDS is also non-alternating. This is discussed in the next section before the completeness and soundness proofs.

Finally, we remark the algorithm is PTIME for a fixed order and number of control states. If we obtained \mathcal{C} from a higher-order recursion scheme, the number of control states is given by the arity of the scheme [10] and the size of the property automaton (giving \mathcal{L}_{err}). In practice, we expect the arity and order to be small, and since simple reachability properties require small automata, we expect the total number of control states to be small.

5.1. Non-Alternation. We introduce a notion of non-alternation at order- n . Note that the automata are alternating both via transitions to Q with $|Q| > 1$, and via collapse links.

Definition 5.3 (Non-Alternation at Order- n). An order- n stack automaton A is *non-alternating at order- n* whenever, for all stacks $w \in \mathcal{L}_q(A)$ for some state q of A , there is an accepting run of A over the graph of w such that no node is labelled by $Q \subseteq \mathbb{Q}_n$ and $|Q| > 1$.

For example, take a run over an example order-2 collapsible stack



This is non-alternating at order- n whenever $|Q_1^1| \leq 1$ and $|Q_2^2| \leq 1$.

Note, for example, that a stack automaton that does not follow collapse links, and has no alternating transitions in Δ_n , is trivially non-alternating at order- n . Similarly, we may allow $q \xrightarrow[Q_{col}]{a} (\emptyset, \dots, \emptyset)$ when $Q_{col} \subseteq \mathbb{Q}_n$ and $|Q_{col}| \leq 1$.

We then define Π' to be the saturation function Π with the additional constraint that a transition $q \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ is not added if $|Q_n| > 1$. Clearly saturation by Π' remains sound, since it contains a subset of the transitions of the automaton produced by saturation with Π . Hence, we only need to prove that the automaton remains complete. We prove completeness in conjunction with the completeness proof for the saturation algorithm in general. Intuitively, the automaton remains correct because a collapse link at order- n can only be used once, whereas, at lower orders, a $push_k$ operation may make different copies of a link (with different targets). Hence, lower order links need alternation to keep track of the different uses of the link throughout the run. Note that restricting the transitions added does not cause any soundness violations.

Henceforth, we will refer to non-alternation at order- n as simply *non-alternation*.

5.2. Completeness. We show that the automaton constructed by Π (and Π') is complete for collapsible stacks. The intuition behind the completeness proof is well illustrated by the examples in Section 4.3, hence we encourage the reader to consult these examples when reading the proof.

Lemma 5.4 (Completeness of Π). *Given a CPDS \mathcal{C} and an order- n stack automaton A_0 , the automaton A constructed by saturation with Π is such that $\langle p, w \rangle \in Pre_{\mathcal{C}}^*(A_0)$ implies $w \in \mathcal{L}_{q_p}(A)$. The result also holds for Π' when \mathcal{C} and A_0 are non-alternating.*

Proof. The proof is by induction over α such that $\langle p, w \rangle \in \text{Pre}_{\mathcal{C}}^{\alpha}(A_0)$. We prove simultaneously during the induction that in the case of Π' all $\langle p, w \rangle \in \text{Pre}_{\mathcal{C}}^*(A_0)$ have an accepting run from q_p of A over the graph of w such that no node is labelled by $Q \subseteq \mathbb{Q}_n$ and $|Q| > 1$.

In the base case, we have $w \in \mathcal{L}_{q_p}(A_0)$ and the existence of a (non-alternating) run of A_0 , and thus a run in A comes directly from the (non-alternating) run of A_0 .

Hence, inductively assume $\langle p, w \rangle \longrightarrow \langle p', w' \rangle$ and a (non-alternating) accepting run of w' from $q_{p'}$ of A via a rule $(p, \text{top}_1(w), o, p')$. Hence $w' = o(w)$.

- (1) When $o = \text{pop}_k$, let $q_{p'} \xrightarrow{q_k} (Q_1, \dots, Q_{k+1})$ be the order- n to order- k part of the initial transition accepting w' . We know, from the construction, we have the transition $q_p \xrightarrow[\emptyset]{a} (\emptyset, \dots, \emptyset, \{q_k\}, Q_{k+1}, \dots, Q_n)$. Let $w = u_{k-1} :_k u_k :_{(k+1)} \dots :_n u_n$. We know $w' = u_k :_{(k+1)} \dots :_n u_n$ and the run from q_k over w' is accepting (and non-alternating). Hence, via the transition from q_p we get an accepting run of w which labels all nodes in w' with the same nodes as the accepting run of w' from q_k . For Π' , since the transitions reading u_{k-1} are all to the empty-set, this gives us a run over w that remains non-alternating.
- (2) When $o = \text{push}_k$, let $w = u_{k-1} :_k \dots :_n u_n$. We know

$$w' = u_{k-1} :_k u_{k-1} :_k u_k :_{(k+1)} \dots :_n u_n .$$

Let $q_{p'} \xrightarrow[\text{Q}_{col}]{a} (Q_1, \dots, Q_k, \dots, Q_n)$ and $Q_k \xrightarrow[\text{Q}'_{col}]{a} (Q'_1, \dots, Q'_k)$ be the initial transitions used on the run of w' (where the transition from Q_k reads the second copy of u_{k-1}).

From the construction we the have a transition

$$q_p \xrightarrow[\text{Q}_{col} \cup \text{Q}'_{col}]{a} (Q_1 \cup Q'_1, \dots, Q_{k-1} \cup Q'_{k-1}, Q'_k, Q_{k+1}, \dots, Q_n) .$$

Since we know $u_k :_{(k+1)} \dots :_n u_n$ is accepted from Q'_k via Q_{k+1}, \dots, Q_n , and we know that u_{k-1} is accepted from Q_1, \dots, Q_{k-1} and Q'_1, \dots, Q'_{k-1} via a -transitions with collapse links Q_{col} and Q'_{col} respectively, we obtain an accepting run of w .

For Π' , since we know that the destinations of the collapse links that go outside of u_{k-1} (in particular, the order- n links) always point to the same stacks in both copies of u_{k-1} in the non-alternating accepting run of w' , we know that the run thus defined from q_p is both accepting and non-alternating. We remark that, while this is *always* the case for order- n links, it is not always the case for order- k links with $k < n$, and hence, alternation is needed at orders lower than n .

- (3) When $o = \text{collapse}_k$, let

$$w = u_{k-1}^1 :_k \dots :_k u_{k-1}^{\ell} :_k u_k :_{(k+1)} \dots :_n u_n$$

where $w' = u_k :_{(k+1)} \dots :_n u_n$ is the stack pointed to by the order- k collapse link of $\text{top}_1(w)$. Let $q_{p'} \xrightarrow{q_k} (Q_{k+1}, \dots, Q_n)$ in A be the initial transition on the accepting run of w' . We know from the construction that we have the transition $q_p \xrightarrow[\{q_k\}]{a} (\emptyset, \dots, \emptyset, Q_{k+1}, \dots, Q_n)$.

We know the run from q_k via Q_{k+1}, \dots, Q_n over w' is accepting (and non-alternating). Hence, via the transition from q_p we get an accepting run of w which labels all nodes in w' with the same nodes as the accepting run of w' from q_k (by following the collapse link). Since the transition reading u_{k-1}^1 is to the empty set, this immediately allows us to label $u_{k-1}^2, \dots, u_{k-1}^{\ell}$ with empty sets, giving us a run over w (that remains non-alternating).

- (4) When $o = push_b^k$, let $w = u_{k-1} :_k u_k :_{k+1} \cdots :_n u_n$. For the appropriate ℓ , we know $w' = push_b^k(w)$ is

$$b^{(k,\ell)} :_1 u_{k-1} :_k \cdots :_n u_n .$$

Let $q_{p'} \xrightarrow[Q_{col}]{b} (Q_1, \dots, Q_n)$ and $Q_1 \xrightarrow[Q'_{col}]{a} Q'_1$ be the first transitions used on the accepting (non-alternating) run of w' . From the construction we know we have $q_p \xrightarrow[Q'_{col}]{a} (Q'_1, Q_2, \dots, Q_k \cup Q_{col}, \dots, Q_n)$ from which we can construct an accepting run of w (which is w' without the first b on top of the top order-1 stack).

For Π' , to see that the run is non-alternating, we observe that $Q_k \cup Q_{col}$ results in the same labelling on the nodes of the graph of w as in corresponding run over w' . Apart from the labelling of the initial nodes (which label the beginning of each top_k stack), the run over w is identical to the run over w' , and hence, non-alternating.

- (5) When $o = rew_b$ let $q_{p'} \xrightarrow[Q_{col}]{b} (Q_1, \dots, Q_n)$ be the first transition on the (non-alternating) accepting run of $w' = b :_1 v$ for some v . From the construction we know we have a transition $q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$, from which we get an accepting, (non-alternating) run of $w = a :_1 v$ as required.

Hence, for every $\langle p, w \rangle \in Pre_C^*(A_0)$ we have $w \in \mathcal{L}_{q_p}(A)$, and when we use Π' and \mathcal{C} and A_0 are non-alternating, the run is non-alternating.

In the alternating case we may have a branching transition $\langle p, w \rangle \longrightarrow C$ (where C is a set of configurations) via a rule $p \rightarrow P$. In this case, let $Q = \{q_{p'} \mid p' \in P\}$ and $Q \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ be the accumulation of all initial transitions from states in Q . By construction, we have a rule $q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ from which an accepting run over w can easily be built. Hence, for every $\langle p, w \rangle \in Pre_C^*(A_0)$ we have $w \in \mathcal{L}_{q_p}(A)$. \square

5.3. Soundness by Witness Generation. In this section, we describe an algorithm that given a CPDS \mathcal{C} and a stack automaton A_0 such that a configuration $\langle p, w \rangle$ of \mathcal{C} belongs to $Pre_C^*(A_0)$, constructs a tree of rules of \mathcal{C} such that, each branch, when applied from $\langle p, w \rangle$ leads to a configuration in $\mathcal{L}(A_0)$. When an alternating rule $p \rightarrow P$ is applied, the tree has a child for each control state appearing in P . Otherwise, each node has a single child.

In practice, we use the algorithm with A_0 accepting the set of all configurations starting with some error state p_{err} . The output is a counter-example showing how the CPDS can reach this error state.

The algorithm is a natural one and the full details are given in the sequel. We describe it informally here by means of the example in Figure 3, described in Section 4.3. In this case, since there is no alternation, we construct a single-branch tree. I.e., a trace.

To construct a trace from $\langle p_1, [b][c][d] \rangle$ to $\langle p_5, [d] \rangle$ we first note that, when adding the initial transition of the pictured run from q_{p_1} , the saturation step marked that the transition was added due to the rule $(p_1, b, push_a^2, p_2)$. If we apply this rule to $\langle p_1, [b][c][d] \rangle$ we obtain $\langle p_2, [ab][c][d] \rangle$ (collapse links omitted). Furthermore, the justifications added during the saturation step tell us which transitions to use to construct the pictured run from q_{p_2} . Hence, we have completed the first step of counter example extraction and moved one step closer to the target configuration. To continue, we consider the initial transition of the run from q_{p_2} . Again, the justifications added during saturation tell us which CPDS rule

to apply and which stack automaton transitions to use to build an accepting run of the next configuration. Thus, we follow the justifications back to a run of A_0 , constructing a complete trace on the way.

The main technical difficulty lies in proving that the reasoning outlined above leads to a terminating algorithm. For example, we need to prove that following the justifications does not result us following a loop indefinitely. Since the stack may shrink and grow during a run, this is a non-trivial property. To prove it, we require a subtle relation on runs over higher-order collapsible stacks.

5.3.1. A Well-Founded Relation on Stack Automaton Runs. We define a well-founded relation over runs of the stack automaton A constructed by saturation from \mathcal{C} and A_0 . To do this we represent a run over a stack as another stack of (sets of) transitions of A . This can be obtained by replacing each instance of a stack character with the set of order-1 transitions that read it. This is formally defined in Section 5.3.2 and described by example here. Consider the run over $[[b][c][d]]$ from q_{p_1} in Figure 3. We can represent this run as the stack $[[\{t_1\}][\{t_2\}][\{t_3\}]]$ where $t_1 = q_5 \xrightarrow{b} \emptyset$, $t_2 = q_2 \xrightarrow{c} \emptyset$ and $t_3 = q_1 \xrightarrow{d} \emptyset$. Note that since q_5 uniquely labels the order-2 transition $q_{p_1} \xrightarrow{q_5} \{q_{p_4}\}$ (and similarly for the transitions from q_{p_4} and q_{p_5}) we do not explicitly store these transitions in our stack representation of runs.

Using this representation, we can define by induction a relation \hookrightarrow_k on the order- k runs of A . Note that this is not an order relation as it is not always transitive. There are several cases to \hookrightarrow_k .

- (1) For $k = 1$, we say $w' \hookrightarrow_1 w$ if for some $i \geq 0$, w contains strictly fewer transitions in Δ_1 justified at step i than w' and that for all $j > i$ they both contain the same number of transitions in Δ_1 justified at step j .
- (2) For $k > 1$, we say $u = [u_\ell \dots u_1]_k \hookrightarrow_k v = [v_{\ell'} \dots v_1]_k$ if
 - (a) $\ell' < \ell$ and $u_i = v_i$ for $i \in [1, \ell' - 1]$ and either $u_{\ell'} = v_{\ell'}$ or $u_{\ell'} \hookrightarrow_{k-1} v_{\ell'}$, or
 - (b) $\ell' \geq \ell$ and $u_i = v_i$ for $i \in [1, \ell - 1]$ and $u_\ell \hookrightarrow_{k-1} v_\ell$ for all $i \in [\ell, \ell']$.

The proof of the following lemma is given in Section 5.3.5.

Lemma 5.5. *For all $k \in [1, n]$, the relation \hookrightarrow_k is well-founded. Namely there is no infinite sequence $w_0 \hookrightarrow_k w_1 \hookrightarrow_k w_2 \hookrightarrow_k \dots$.* \square

It is possible to show that by following the justifications, from stack w to a w' , we always have $w \hookrightarrow_n w'$. Since this relation is well-founded, witness generation always terminates.

5.3.2. Alternative definition of the runs. We give an alternative definition of a run of a stack automaton that is more appropriate to perform the *run surgery* needed below. The definition of an accepting run requires two intermediary notions.

A run of A on an order- $\leq n$ stack v is an annotation of each symbol of this stack by a subset of Δ_1 , the order-1 transitions of A . Formally a run over an order- k stack v is an order- k stack over the alphabet $\Sigma \times 2^{\Delta_1}$ such that when projecting on the Σ -component we retrieve the stack v .

Let w be an order- k run of A . For a set $Q \subseteq \mathbb{Q}_k$ of order- k states, we say that w is *Q -valid* if the following holds. If the run w is empty, Q must be a subset of \mathcal{F}_k . Assume now that w is not empty. If $k = 1$ and $w = (a, T) :_1 w'$, there must exist $Q' \subseteq \mathbb{Q}_1$ such that

w' is Q' -valid and for all $q \in Q$, there exists a transition in T of the form $q \xrightarrow[Q_{col}]{a} Q''$ with $Q'' \subseteq Q'$. If $k > 1$ and $w = u :_k w'$ then there must exist a subset $Q' \subseteq Q_k$ of order- k states such that w' is Q' -valid and for all $q \in Q$, there exists a transition $q \xrightarrow{q_{Q''}} Q'' \in \Delta_k$ such that u is $\{q_{Q''}\}$ -valid.

Note that Q -validity does not check the constraint imposed by the Q_{col} component appearing in order-1 transitions. This is done by *link-validity* which is only meaningful on order- n runs: An order- n run w is *link-valid* if for every substack of w of the form $w' = (a, T)^{k,i} : w''$ and every transition $q \xrightarrow[Q_{col}]{a} Q$ then $top_{k+1}(collapse_k(w'))$ is Q_{col} -valid.

For $q \in Q_n$, an order- n run w is q -accepting if it is both $\{q\}$ -valid and link-valid. In addition, we require that if w is non empty and hence of the form $(a, T) :_1 w'$ then T is reduced to a singleton $\{t\}$ and we refer to t as the head transition of the run.

5.3.3. Witness Trees. We define what it means to be a witness of $\langle p, w \rangle \in Pre_C^*(A_0)$. Without alternation, we simply require a trace of rules which take $\langle p, w \rangle$ to some configuration in $\mathcal{L}(A_0)$. In the presence of alternating transitions $p \rightarrow P$ we need to account for each possible next configuration. Hence, we require finite trees rather than sequences.

A Γ -labelled *finite tree* is a tuple (D, λ) where $D \subset \mathbb{N}^*$ is a tree domain that is both prefix- and younger-sibling-closed. That is, for all $\eta i \in D$ with $\eta \in \mathbb{N}^*$ and $i \in \mathbb{N}$ we have $\eta \in D$ and moreover, for all $j < i$ we have $\eta j \in D$. Furthermore, $\lambda : D \rightarrow \Gamma$ is a tree labelling for a set Γ of labels. A leaf node is a node $\eta \in D$ such that there is no i with $\eta i \in D$. Otherwise, η is an internal node.

Definition 5.6 (Witness Trees). For a CPDS \mathcal{C} , configuration $\langle p, w \rangle$ and stack automaton A_0 , a witness tree is a Γ -labelled finite tree (D, λ) where Γ contains labels of the form $[c]$ and $[c, r]$ with c a configuration of \mathcal{C} and r a rule of \mathcal{C} . Moreover, for all $\eta \in D$ we have

- if $\eta = \varepsilon$ then $\lambda(\eta) = [c]$ or $\lambda(\eta) = [c, r]$ for some r and $c = \langle p, w \rangle$, and
- if η is an internal node then $\lambda(\eta) = [\langle p_1, w_1 \rangle, r]$ for some p_1, w_1 , and r , and
 - if $r = (p_1, a, o, p_2)$ then $\eta 1$ is the only child of η and $\lambda(\eta 1) = [\langle p_2, w_2 \rangle]$ or $\lambda(\eta 1) = [\langle p_2, w_2 \rangle, r']$ with $w_2 = o(w_1)$, and
 - if $r = (p_1, P)$ then for each $p_2 \in P$ there is some $\eta i \in D$ such that $\lambda(\eta i) = [\langle p_2, w_1 \rangle]$ or $\lambda(\eta i) = [\langle p_2, w_1 \rangle, r']$, and
- if η is a leaf node then $\lambda(\eta) = [c]$ for some $c \in \mathcal{L}(A_0)$.

The following proposition gives us the required property of witness trees that allows us to use them to prove soundness.

Proposition 5.7. *For a CPDS \mathcal{C} , stack automaton A_0 and configuration c of \mathcal{C} , if there is a witness tree for c , then $c \in Pre_C^*(A_0)$.*

Proof. A straightforward induction beginning at the leaves of the witness tree. □

Algorithm 1 Counter-example Extraction

Require: A stack-automaton A generated by saturating A_0 .

Ensure: Print a finite sequence of CPDS rules that when executed will lead from a configuration $\langle p_0, w \rangle \in \mathcal{L}(A)$ to one in $\mathcal{L}(A_0)$.

Fix w to be a trimmed accepting run of w from q_{p_0}

return GetWitness(w)

Algorithm 2 GetWitness(\mathbf{w})

Require: A trimmed accepting run of a stack u from a control state p .**Ensure:** A witness tree for $\langle p, u \rangle$.

```

if the head transition  $t$  of  $\mathbf{w}$  is justified by 0 then
  return  $[\langle p, u \rangle]$ 
else
  Let  $r$  be the CPDS rule appearing in the justification of  $t$ .
  Let  $\alpha = [\langle p, u \rangle, r]$ .
  if  $r = (p, a, pop_k, p')$  for some  $1 \leq k \leq n$  then
    The transition  $t$  is of the form  $q_p \xrightarrow[a]{a} (\emptyset, \dots, \emptyset, \{q_{p', Q_n, \dots, Q_{k+1}}\}, Q_{k+1}, \dots, Q_n)$ 
    Pick  $Q_k, \dots, Q_1, Q_{col}$  such that  $t' := q_{p'} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_k, Q_{k+1}, \dots, Q_n)$  is in
     $top_1(pop_k(\mathbf{w}))$ 
    return  $\alpha(\text{GetWitness}(\text{rew}_{\{t'\}}(pop_k(\mathbf{w}))))$ 
  else if  $r = (p, a, collapse_k, p')$  for  $2 \leq k \leq n$  then
    The transition  $t$  is of the form  $t = q_p \xrightarrow[\{q_{p', Q_n, \dots, Q_{k+1}}\}]{a} (\emptyset, \dots, \emptyset, Q_{k+1}, \dots, Q_n)$ 
    Pick  $Q_k, \dots, Q_1, Q_{col}$  such that  $t' := q_{p'} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_k, Q_{k+1}, \dots, Q_n)$  is in
     $top_1(collapse_k(\mathbf{w}))$ 
    return  $\alpha(\text{GetWitness}(\text{rew}_{\{t'\}}(collapse_k(\mathbf{w}))))$ 
  else if  $r = (p, a, rew_b, p')$  for some  $b \in \Sigma$  then
     $J(t)$  must be of form  $(r, t', i)$ 
    return  $\alpha(\text{GetWitness}(\text{rew}_{\{t'\}}(rew_b(\mathbf{w}))))$ 
  else if  $r = (p, a, o, p')$  with  $o = push_k$  or  $o = push_b^{k'}$  then
     $J(t)$  must be of the form  $(r, t', T, i)$ 
    return  $\alpha(\text{GetWitness}(\text{rew}_{\{t'\}}(o(\text{rew}_T(\mathbf{w})))))$ 
  else if  $r = p \rightarrow P$  then
     $J(t)$  must be of the form  $(r, T, i)$ 
     $P$  is of the form  $\{p_1, \dots, p_\ell\}$ 
    For each  $j$  we have a transition of the form  $t_j := q_{p_j} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$  in  $T$ .
    return  $\alpha(\text{GetWitness}(\text{rew}_{\{t_1\}}(\mathbf{w})), \dots, \text{GetWitness}(\text{rew}_{\{t_\ell\}}(\mathbf{w})))$ 

```

5.3.4. *The Algorithm.* Algorithm 1 and Algorithm 2 shows how we construct counter examples. Variable \mathbf{w} contains a run of A which is q_p -accepting for some state p . The initial value of \mathbf{w} , denoted w_0 , is an accepting run for the initial configuration $\langle p_0, u_0 \rangle$. We construct \mathbf{w} recursively, with each recursive call building a different branch of the witness tree. At the beginning of each recursive call, let w be the value of \mathbf{w} which we assume to be an accepting run for a configuration $\langle p, u \rangle$. Moreover, let t denote the head transition of w .

Each recursive call returns a witness tree from $\langle p, u \rangle$. Moreover and crucially for termination, each recursive call with argument value w' is such that $w \hookrightarrow_n w'$. As \hookrightarrow_n is well-founded, the recursive calls eventually reach the base case after N calls with the final t justified by 0. It is then possible to prune the last run w to form a run that consists entirely of transitions already belonging to A_0 by the assumption that initial states at every level

of A have no incoming transitions. It follows that the configurations reached at the leaves of the witness tree belongs to $\mathcal{L}(A_0)$.

5.3.5. Correctness of the Algorithm. In this section, we establish the correctness of Algorithm 1 and give omitted proofs. We start with the proof of Lemma 5.5.

Proof. For $k = 1$, consider for any order-1 run w the tuple $|w| = (n_m, \dots, n_0)$ where m is the step at which the saturation algorithm terminates and for all $i \in [0, m]$, n_i is the number of occurrences in w of transitions in Δ_1 justified at step i . The relation \hookrightarrow_1 can be equivalently defined as $w \hookrightarrow_1 w'$ if $|w'|$ is lexicographically smaller than $|w|$. It immediately follows that \hookrightarrow_1 is well-founded.

For $k + 1 > 1$ assuming the property holds for \hookrightarrow_k . Suppose for contradiction that \hookrightarrow_{k+1} is not well-founded. Then there must be an infinite chain of runs of the form:

$$w_1 \hookrightarrow_{k+1} w_2 \hookrightarrow_{k+1} w_3 \hookrightarrow_{k+1} \dots$$

Now pick an index i such that for every $j > i$ it is the case that w_j is at least as long (w.r.t the number of order- $(k-1)$ stacks) as the run w_i (infinitely many such indices must clearly exist since comparing runs by their lengths is a well-founded relation.). If $w_i = u :_{k+1} w'_i$, it is a straightforward induction to see that for every $j > i$ w_j is of the form $w''_j w'_j$ with $u \hookrightarrow_k^+ v$ for all order- k run v occurring in w''_j where \hookrightarrow_k^+ designates the transitive closure of \hookrightarrow_k .

So in particular if we pick infinitely many positions in the chain i_ℓ such that the run $w_{i_\ell} = u_{i_\ell} :_{k+1} w'_{i_\ell}$ is at least as long as the sequence w_j for all $j > i_\ell$ it must be the case that:

$$u_{i_1} \hookrightarrow_k^+ u_{i_2} \hookrightarrow_k^+ u_{i_3} \hookrightarrow_k^+ \dots$$

This in turn contradicts the fact that \hookrightarrow_k is well-founded. \square

The next lemma describes two sufficient conditions condition for $w \hookrightarrow_k w'$ to hold.

Lemma 5.8. *The following properties hold:*

- (1) *Let w and w' be two order- n runs such that for some $1 \leq k < n$, $\text{top}_{k+1}(w) \hookrightarrow_k \text{top}_{k+1}(w')$ and $\text{pop}_{k+1}(w) = \text{pop}_{k+1}(w')$ then $w \hookrightarrow_n w'$.*
- (2) *Let w be an order- k run and let T be a set of transitions that is smaller than some transition appearing in $\text{top}_1(w)$, $w \hookrightarrow_k \text{rew}_T(w)$.*

Proof. For the first property, we will show by induction on k' that for all $k' \in [k, n]$, $\text{top}_{k'+1}(w) \hookrightarrow_{k'} \text{top}_{k'+1}(w')$. The case $k' = k$ is assumed to hold in the hypothesis. Assume that the property holds for k' . We have $\text{top}_{k'+2}(w) = \text{top}_{k'+1}(w) :_{k'} \text{top}_{k'+2}(\text{pop}_{k'+1}(w))$ and $\text{top}_{k'+2}(w') = \text{top}_{k'+1}(w') :_{k'} \text{top}_{k'+2}(\text{pop}_{k'+1}(w'))$. Remark that $\text{pop}_{k'+1}(w) = \text{pop}_{k'+1}(w') = u$. This is by assumption for $k = k'$ and if $k' > k$ $\text{pop}_{k'+2}(w) = \text{pop}_{k'+2}(\text{pop}_{k+1}(w)) = \text{pop}_{k'+2}(\text{pop}_{k+1}(w')) = \text{pop}_{k'+2}(\text{pop}_{k+1}(w'))$. Hence $\text{top}_{k'+2}(w) = \text{top}_{k'+1}(w) :_{k'} u$ and $\text{top}_{k'+2}(w') = \text{top}_{k'+1}(w') :_{k'} u$ with $\text{top}_{k'+1} \hookrightarrow_{k'} \text{top}_{k'+1}(w')$. By definition of $\hookrightarrow_{k'+1}$, we have $\text{top}_{k'+2} \hookrightarrow_{k'+1} \text{top}_{k'+2}(w')$.

For the second property, we have $\text{top}_2(w) \hookrightarrow_1 \text{top}_2(\text{rew}_T(w))$ (by definition of \hookrightarrow_1) and $\text{pop}_2(w) = \text{pop}_2(w')$. Hence by the first-property $w \hookrightarrow_k \text{rew}_T(w)$. \square

We now prove Algorithm 1 correct. As it is often the case, we restrict our attention to runs containing only “useful” transitions. A run w is *trimmed* if for any o_1, \dots, o_j of *pop* operations producing a subrun $w' = o_j(\dots o_1(w) \dots)$, for any order-1 transition

$$q \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$$

appearing in $top_1(w')$, we have for all $i \in [1, m-1]$, that $top_{i+1}(pop_i(w'))$ is Q_i -valid where m is the smallest index such that pop_m appears in the sequence o_1, \dots, o_j .

Proposition 5.9. *Algorithm 1 is correct.*

Proof. The initial value of w , denoted w_0 , is an accepting run for the initial configuration $\langle p_0, u_0 \rangle$. An updated value of w is passed at each recursive call. We denote by w the value of w at the beginning of each call.

We are going to prove by induction on the depth of recursion that w is always a trimmed q_p -accepting run on some stack u . Furthermore, for each recursive call we have $w \hookrightarrow_n w'$ where w' is the value passed to the call.

For $i = 0$, w_0 which contains only one transition is necessarily trimmed. Assume that the property holds for w , and let us prove it for each w' appearing in a recursive call. By the induction hypothesis, w is a trimmed q_p -accepting run on a stack u' . This implies that its head transition is of the form:

$$t = q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n).$$

Hence its justification contains a transition of the CPDS of the form (p, a, o, p') or $p \rightarrow P$. In the first case, we reason by case distinction on the operation o .

If $o = rew_b$ for some $b \in \Sigma$. The transition t is of the form $q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ with a justification of the form $J(t) = (r, t', i)$ with t' of the form $q_{p'} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$. Note that t' was introduced before t .

The run w' is equal to $rew_{\{t'\}}(rew_b(w))$. It is clear that w' is a trimmed $q_{p'}$ -accepting run on the stack $rew_b(u)$. By the second property of Lemma 5.8, $w \hookrightarrow_n w'$.

If $o = pop_k$ for some $k \in [1, n]$. The transition t is of the form

$$q_p \xrightarrow[\emptyset]{a} (\emptyset, \dots, \emptyset, \{q_{p', Q_n, \dots, Q_{k+1}}\}, Q_{k+1}, \dots, Q_n).$$

As w is q_p -accepting, it follows that for all $j \in [k+1, n]$, $top_{j+1}(pop_j(w))$ is Q_j -valid and that $top_{k+1}(pop_k(w))$ is $\{q_*\}$ -valid for $q_* = q_{p', Q_n, \dots, Q_{k+1}}$. By unfolding the notion of $\{q_*\}$ -validity, we obtain that $top_1(pop_k(w))$ contains at least one transition t' of the form:

$$q_{p'} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_k, Q_{k+1}, \dots, Q_n)$$

Let t' be the transition of this form picked by the algorithm. As w is trimmed it follows that for all $j \in [1, k]$, $top_{j+1}(pop_j(pop_k(w)))$ is Q_j -valid.

We have $w' = rew_{\{t'\}}(pop_k(w))$. As w' is a subrun of w (which is link-valid and trimmed), it is link-valid and trimmed. To prove that it is also $q_{p'}$ -valid it is enough to show that for all $i \in [1, n]$, we have $top_{i+1}(pop_i(w'))$ is Q_i -valid. For $i \in [k+1, n]$, we have seen that $top_{i+1}(pop_i(w')) = top_{i+1}(pop_i(w))$ is Q_i -valid. For $i \in [1, k]$, we have seen that $top_{i+1}(pop_i(w')) = top_{i+1}(pop_i(pop_k(w)))$ is Q_i -valid.

It only remains to show that $w \hookrightarrow_n w'$. By the first property of Lemma 5.8, it is enough to show that $top_{k+1}(w) \hookrightarrow_k top_k(w')$ (as $pop_{k+1}(w') = pop_{k+1}(w)$ if $k < n$). First consider

the case when $k = 1$. It follows from the fact that the set of order-1 transitions appearing in $top_2(w')$ is strictly included in $top_2(w)$. Now assume that $k > 1$. The run $top_{k+1}(w)$ can be written as $u :_{k+1} u' :_{k+1} v$ and $top_{k+1}(w) = rew_T(u') :_{k+1} v$. By the second property of Lemma 5.8, $u' \hookrightarrow_{k-1} rew_T(u')$ and by definition of \hookrightarrow_k , $top_{k+1}(w) \hookrightarrow_k top_k(w')$.

If $o = collapse_k$ for some $k \in [2, n]$. This case is similar to the pop_k case.

If $o = push_k$ for some $k \in [2, n]$. The transition t is of the form

$$t = q_p \xrightarrow[Q_{col} \cup Q'_{col}]^a (Q_1 \cup Q'_1, \dots, Q_{k-1} \cup Q'_{k-1}, Q'_k, Q_{k+1}, \dots, Q_n)$$

with $J(t) = (r, t', T, i + 1)$ where

$$t' = q_{p'} \xrightarrow[Q_{col}]^a (Q_1, \dots, Q_k, \dots, Q_n)$$

and T is a set of transitions of the form $Q_k \xrightarrow[Q'_{col}]^a (Q'_1, \dots, Q'_k)$.

The run w' is equal to $rew_{\{t'\}}(o(rew_T(w)))$. Let $w = u :_k v$. The run w' is then equal to $rew_{\{t'\}}(u) :_k rew_T(u) :_k v$.

Let us first show that w' is $\{q_p\}$ -valid. For this it is enough to show that:

- for all $k' \in [k + 1, n]$, $top_{k'+1}(pop_{k'}(w')) = top_{k'+1}(pop_{k'}(w))$ is $Q_{k'}$ -valid. This immediately follows from the fact that w is q -accepting with head transition t .
- $top_k(pop_k(w')) = rew_T(u) :_k v = rew_T(w)$ is Q_k -valid. As T has the form $Q_k \xrightarrow[Q'_{col}]^a (Q'_1, \dots, Q'_k)$, it enough to show that for all $k' \in [1, k]$, we have $top_{k'+1}(pop_{k'}(rew_T(w))) = top_{k'+1}(pop_{k'}(w))$ is $Q_{k'}$ -valid. This immediately follows from the fact that w is q_p -accepting with head transition t .
- for all $k' \in [1, k - 1]$, $top_{k'+1}(pop_{k'}(w')) = top_{k'+1}(pop_{k'}(w))$ is $Q_{k'}$ -valid. This immediately follows from the fact that w is q -accepting with head transition t .

We now show that w' is link-valid. We only need to check the validity for the substack $rew_T(u) :_k v$ and the substacks of the form $u' :_k rew_T(u) :_k v$ where u' is a substack of $rew_{\{t'\}}(u)$. Let us first consider the stack $rew_T(u) :_k v$ and let h be a transition in T of the form

$$q_h \xrightarrow[Q_{col}^h]^a (Q_1^h, \dots, Q_n^h)$$

We have that Q_{col}^h is a subset of Q'_{col} . Let k' be the order of the link on top of $rew_T(u) :_k v$. As w is link-valid, we now that $top_{k'+1}(collapse_{k'}(w)) = top_{k+1}(collapse_{k'}(rew_T(u) :_k v))$ is $Q_{col} \cup Q'_{col}$ -valid hence it is also Q_{col}^h -valid. We now move on to the case of $x = rew_{\{t'\}}(u) :_k rew_T(u) :_k v$. Let k' be the order of the link on top of x . We have that $top_{k'+1}(collapse_{k'}(x)) = top_{k'+1}(collapse_{k'}(w))$. By link-validity of w , it is the case that $top_{k'+1}(collapse_{k'}(w))$ is $Q_{col} \cup Q'_{col}$ -valid and in particular Q_{col}^h -valid.

Finally let u' be a strict substack of $rew_{\{t'\}}(u)$. Let k' be the order of the link appearing on top of $x = u' :_k rew_T(u) :_k v$ and let h be a transition attached to the top of x of the form:

$$q_h \xrightarrow[Q_{col}^h]^a (Q_1^h, \dots, Q_n^h)$$

We have that $top_{k'+1}(collapse_{k'}(x)) = top_{k'+1}(collapse_{k'}(w))$. By link-validity of w , it is the case that $top_{k'+1}(collapse_{k'}(w))$ is Q_{col}^h -valid.

It now remains to show that w' is trimmed. The only interesting case is that of the substack $rew_T(u) :_k v$ which is reach by a pop_k operation. Any transition $h \in T$, is of the

form

$$q_h \xrightarrow[Q_{col}^h]{a} (Q_1^h, \dots, Q_n^h)$$

with for all $k' \in [1, k]$, $Q_{k'}^h \subseteq Q_{k'}'$. Hence it is enough for us to show that for all $k' \in [1, k]$, $top_{k'+1}(pop_{k'}(rew_T(u) :_k v)) = top_{k'+1}(pop_{k'}(w))$ is $Q_{k'}'$ -valid. This immediately follows from the fact that w is q -accepting with head transition t .

It only remains to show that $w \hookrightarrow_n w'$. First remark that $u \hookrightarrow_{k-1} rew_{\{t'\}}(u)$ and $u \hookrightarrow_{k-1} rew_T(u)$ as in both cases t is replaced by one or several transition with a smaller timestamp (cf. second property of Lemma 5.8). By definition of \hookrightarrow_k , we have $top_{k+1}(w) \hookrightarrow_k top_{k+1}(w')$. The first property of Lemma 5.8 then implies that $w \hookrightarrow_k w'$.

If $o = push_b^k$ for some $b \in \Sigma$ and $k \in [2, n]$. This case is similar to the $push_k$ case.

This concludes the case where the justification contains a transition of the form (p, a, o, p') .

When t is of the form $p \rightarrow P$ then the transition t is of the form $q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$

with a justification of the form $J(t) = (r, T, i)$ with T of the form $Q \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ for

$Q = \{q_{p'} \mid p' \in P\}$. Thus for each $p' \in P$ we have some transition $t' = q_{p'} \xrightarrow[Q_{col}']{a} Q_1', \dots, Q_n'$.

Note that t' was introduced before t . The run w' in the corresponding recursive call is equal to $rew_{\{t'\}}(w)$. It is clear that w' is a trimmed $q_{p'}$ -accepting run on the stack u . By the second property of Lemma 5.8, $w \hookrightarrow_n w'$.

In all cases, the recursive call is made with a smaller stack. Since \hookrightarrow_n is well-founded, we eventually reach the base case of the recursion. Thus the algorithm terminates.

That the algorithm returns a witness tree can be proven by induction from the leaves of the recursion back to the beginning of the algorithm. It is immediate in the base case, since a justification of 0 implies that the configuration is accepted by A_0 . When a rule of the form (p, a, o, p') leads to the recursive call we know by induction that we obtain a witness tree for $\langle p', o(u) \rangle$. By adding α as the root of this tree, we immediately get a witness tree for $\langle p, u \rangle$. The remaining case is when a rule $p \rightarrow P$ is used. For each $p' \in P$ we obtain a witness tree for $\langle p', u \rangle$. By constructing the tree with α at the root and children from each of the recursive calls, we have a witness tree for $\langle p, u \rangle$ as required. \square

6. INITIAL FORWARD ANALYSIS

It is generally completely impractical to compute $Pre_C^*(\mathcal{E})$ in full (most non-trivial examples considered in our experiments would time-out). For our saturation algorithm to be usable in practice, it is therefore essential that the search space is restricted, which we achieve by means of an initial forward analysis of the CPDS. In short, we compute an over-approximation of all reachable configurations, and try to restrict our backwards reachability analysis to only include configurations in this over-approximation.

In this section we distinguish an *error state* p_{err} and we are interested only in whether \mathcal{C} can reach a configuration of the form $\langle p_{err}, w \rangle$. That is, the set of target configurations is $\{\langle p_{err}, w \rangle \mid w \in Stacks_n\}$. This suffices to capture the same safety (reachability) properties of recursion schemes as TRecS [19].

We fix a stack-automaton \mathcal{E} recognising all error configurations (those with the state p_{err}). We write $Post_C^*$ for the set of configurations reachable by \mathcal{C} from the initial configuration c_0 . More formally, we have $Post_C^*$ is the smallest set such that $c_0 \in Post_C^*$

and

$$Post_{\mathcal{C}}^* \supseteq \left\{ \langle p', w' \rangle \mid \exists \langle p, w \rangle \in Post_{\mathcal{C}}^* \text{ with } \left(\begin{array}{l} \langle p, w \rangle \longrightarrow \langle p', w' \rangle \vee \\ \langle p, w \rangle \longrightarrow C \text{ and } \langle p', w' \rangle \in C \end{array} \right) \right\}.$$

This set cannot be represented precisely by a stack automaton [2] (for instance using *push*₂, we can create $[[a^n]_1[a^n]_1]_2$ from $[[a^n]_1]_2$ for any $n \geq 0$). We summarise our approach then give details in Sections 6.1, 6.2 and 6.3. Note that the handling of alternating transitions in $Post_{\mathcal{C}}^*$ allows us to treat these transitions in our approximation algorithms in the same way as if the choice were non-deterministic rather than alternating.

Ideally we would compute only $Pre_{\mathcal{C}}^*(\mathcal{E}) \cap Post_{\mathcal{C}}^*$. Since this cannot be represented by an automaton, we instead compute a sufficient approximation T (ideally a *strict* subset of $Pre_{\mathcal{C}}^*(\mathcal{E})$) where:

$$Pre_{\mathcal{C}}^*(\mathcal{E}) \cap Post_{\mathcal{C}}^* \subseteq T \subseteq Pre_{\mathcal{C}}^*(\mathcal{E}).$$

The initial configuration will belong to T iff it can reach a configuration recognised by \mathcal{E} . Computing such a T is much more feasible.

We first compute an over-approximation of $Post_{\mathcal{C}}^*$. For this we use a *summary algorithm* [29] (that happens to be precise at order-1) from which we extract an over-approximation of the set of CPDS rules that may be used on a run to p_{err} . Let \mathcal{C}' be the (smaller) CPDS containing only these rules. I.e., we remove all rules that we know cannot appear on a run to p_{err} . We could thus take $T = Pre_{\mathcal{C}'}^*(\mathcal{E})$ (computable by saturation for \mathcal{C}') since it satisfies the conditions above. This is what we meant by ‘*pruning*’ the CPDS (1a on page 3).

However, we further improve performance by computing an even smaller T (1b in the list on page 3). We extract contextual information from our over-approximation of $Post_{\mathcal{C}}^*$ about how pops and collapses might be used during a run to p_{err} . Our \mathcal{C}' is then restricted to a model \mathcal{C}'' that ‘guards’ its rules by these contextual constraints. Taking $T = Pre_{\mathcal{C}''}^*(\mathcal{E})$ we have a T smaller than $Pre_{\mathcal{C}'}^*(\mathcal{E})$, but still satisfying our sufficient conditions. In fact, \mathcal{C}'' will be a ‘*guarded CPDS*’ (defined in the next subsection). We cannot compute $Pre_{\mathcal{C}''}^*(\mathcal{E})$ precisely for a guarded CPDS, but we can adjust saturation to compute T such that $Pre_{\mathcal{C}''}^*(\mathcal{E}) \subseteq T \subseteq Pre_{\mathcal{C}'}^*(\mathcal{E})$. This set will thus also satisfy our sufficient conditions.

6.1. Guarded Destruction. An *order- n guarded CPDS* (*n -GCPDS*) is an n -CPDS where conventional pop_k and $collapse_k$ operations are replaced by *guarded operations* of the form pop_k^S and $collapse_k^S$ where $S \subseteq \Sigma$. These operations may only be fired if the resulting stack has a member of S on top. That is, for $o \in \{collapse_k, pop_k \mid 1 \leq k \leq n\}$:

$$o^S(u) := \begin{cases} o(u) & \text{if } o(u) \text{ defined and } top_1(o(u)) \in S \\ \text{undefined} & \text{otherwise} \end{cases}.$$

Note, we do not guard the other stack operations since these themselves guarantee the symbol on top of the new stack (e.g. when a transition $(p, a, push_2, p')$ fires it must always result in a stack with a on top, and $(p, a, push_k^b, p')$ produces a stack with b on top).

For a GCPDS \mathcal{C} , we write $\mathbf{Triv}(\mathcal{C})$ for the *trivialisation* of \mathcal{C} : the ordinary CPDS obtained by replacing each pop_k^S (resp. $collapse_k^S$) in the rules of \mathcal{C} with pop_k (resp. $collapse_k$).

We modify the saturation algorithm to use ‘guarded’ saturation steps for pop and collapse rules. Other saturation steps are unchanged. Non-trivial guards reduce the size of the stack-automaton constructed by avoiding additions that are only relevant for unreachable (and hence uninteresting) configurations in the pre-image. Thus, we improve performance.

- (1) when $o = \text{pop}_k^S$, for each $q_{p'} \xrightarrow{q_k} (Q_{k+1}, \dots, Q_n)$ in A such that there is a transition $q_k \xrightarrow{b} (-, \dots, -)$ in A such that $b \in S$, add $q_p \xrightarrow{a} (\emptyset, \dots, \emptyset, \{q_k\}, Q_{k+1}, \dots, Q_n)$ to A' ,
- (3) when $o = \text{collapse}_k^S$, for each $q_{p'} \xrightarrow{q_k} (Q_{k+1}, \dots, Q_n)$ in A where there is a transition $q_k \xrightarrow{b} (-, \dots, -)$ in A with $b \in S$, add $q_p \xrightarrow[\{q_k\}]{a} (\emptyset, \dots, \emptyset, Q_{k+1}, \dots, Q_n)$ to A' .

E.g., suppose that an ordinary (non-guarded) 2-CPDS has rules $(p_1, c, \text{collapse}_2, p)$ and $(p_2, d, \text{collapse}_2, p')$. The *original* saturation algorithm would process these rules to add the transitions: $q_{p_1} \xrightarrow[\{q_p\}]{c} (\emptyset, \emptyset)$ and $q_{p_2} \xrightarrow[\{q_{p'}\}]{d} (\emptyset, \emptyset)$.

Now suppose that the saturation algorithm has produced two transitions of the form $q_p \xrightarrow{a} (-, -)$ and $q_{p'} \xrightarrow{b} (-, -)$. If a GCPDS had, for example, the rules $(p_1, c, \text{collapse}_2^{\{a\}}, p)$ and $(p_2, d, \text{collapse}_2^{\{b\}}, p')$, then these same two transitions would be added by the modified saturation algorithm. On the other hand, the rules $(p_1, c, \text{collapse}_2^{\{a\}}, p)$ and $(p_2, d, \text{collapse}_2^{\{a\}}, p')$ would only result in the first of the two transitions being added.

Lemma 6.1. *The revised saturation algorithm applied to \mathcal{E} (for a GCPDS \mathcal{C}) gives a stack automaton recognising T such that $\text{Pre}_{\mathcal{C}}^*(\mathcal{E}) \subseteq T \subseteq \text{Pre}_{\text{Triv}(\mathcal{C})}^*(\mathcal{E})$* \square

Proof. We can see that $T \subseteq \text{Pre}_{\text{Triv}(\mathcal{C})}^*(\mathcal{E})$ since every time we can add a transition during the modified saturation algorithm we could have added the corresponding guard-free rule in the original, and the original is already known to be sound.

Checking that $\text{Pre}_{\mathcal{C}}^*(\mathcal{E}) \subseteq T$ is an easy modification of the completeness proof for the original algorithm in Lemma 5.4. This works by induction on the length of a path from a configuration in $\text{Pre}_{\mathcal{E}}^*(\mathcal{C})$ to one in \mathcal{E} . Suppose we have a stack-automaton A recognising a configuration (p', u') together with a rule (p, a, o^S, p') of \mathcal{C} where o is either a *pop* or a *collapse* operation. Suppose that (p, u) can reach (p', u') in a single step via this rule. By definition it must then be the case that $\text{top}_1(u') = b$ for some $b \in S$ (and also that $u' = o(u)$). But then the run recognising (p', u') must begin with a transition of the form $q_{p'} \xrightarrow[\text{Q}_{\text{col}}]{b} (Q_1, \dots, Q_n)$. Thus in particular $q_{p', Q_n, \dots, Q_{k+1}} \xrightarrow[\text{Q}_{\text{col}}]{b} (Q_1, \dots, Q_k)$ is the first long-form order- k transition in this run. But then taking $q_k := q_{p', Q_n, \dots, Q_{k+1}}$ we can see that applying the step for the operation o^S in the revised saturation algorithm will create a stack-automaton recognising u . \square

The reason that the algorithm may result in a stack-automaton recognising configurations that do not belong to $\text{Pre}_{\mathcal{C}}^*(\mathcal{E})$ (albeit still in $\text{Pre}_{\text{Triv}(\mathcal{C})}^*(\mathcal{E})$) is that a stack-automaton state q_k emitting a transition $q_k \xrightarrow{b} (-, \dots, -)$ may also emit another transition $q_k \xrightarrow{b'} (-, \dots, -)$ with $b \neq b'$. We could obtain a precise algorithm by taking level- n stack-automaton states of the form $\mathcal{P} \times \Sigma$ so that they represent the top stack-character of a configuration as well as its control-state. However, since Σ is usually large compared to \mathcal{P} and since the worst-case size of the stack-automaton is n -exponential in the number of level- n states this would potentially come at a large practical cost and in any case destroy fixed-parameter tractability. We leave it for future work to investigate how this potential for accuracy could be balanced with the inevitable cost.

Remark 6.2. The above modification to the naive saturation algorithm can also be easily incorporated into the efficient fixed point algorithm described in Section 7.

6.2. Approximate Reachability Graphs. We now describe the summary algorithm used to obtain an over-approximation of $Post_{\mathcal{C}}^*$ and thus compute the GCPDS \mathcal{C}'' mentioned previously. For simplicity, we assume that a stack symbol uniquely determines the order of any link that it emits (which is the case for a CPDS obtained from a HORS). We first describe the approximate reachability graph, and then the approximate summary algorithm.

6.2.1. The Approximate Reachability Graph. We begin with an informal description before the formal definition. Informally, an *approximate reachability graph* for \mathcal{C} is a structure (H, E, B) describing an over-approximation of the reachable configurations of \mathcal{C} .

- The set of nodes of the graph H consists of *heads* of the CPDS, where a head is a pair $(p, a) \in \mathcal{P} \times \Sigma$ and describes configurations of the form $\langle p, u \rangle$ where $top_1(u) = a$.
- The set E contains directed edges $((p, a), r, (p', a'))$ labelled by rules of \mathcal{C} . Such edges over-approximate the transitions that \mathcal{C} might make using a rule r from a configuration described by (p, a) to one described by (p', a') . For example, suppose that \mathcal{C} is order-2 and has, amongst others, the rules $r_1 := (p_1, b, push_2, p_2)$, $r_2 := (p_2, b, push_c^2, p_3)$ and $r_3 := (p_3, c, pop_1, p_4)$ so that it can perform transitions:

$$\left\langle p_1, \left[\begin{array}{c} b \\ a \end{array} \right] \right\rangle \xrightarrow{r_1} \left\langle p_2, \left[\begin{array}{c} b \\ a \end{array} \right] \left[\begin{array}{c} b \\ a \end{array} \right] \right\rangle \xrightarrow{r_2} \left\langle p_3, \left[\begin{array}{c} c \\ b \\ a \end{array} \right] \left[\begin{array}{c} b \\ a \end{array} \right] \right\rangle \xrightarrow{r_3} \left\langle p_4, \left[\begin{array}{c} b \\ a \end{array} \right] \left[\begin{array}{c} b \\ a \end{array} \right] \right\rangle$$

where the first configuration mentioned here is reachable. We should then have edges $((p_1, b), r_1, (p_2, b))$, $((p_2, b), r_2, (p_3, c))$ and $((p_3, c), r_3, (p_4, b))$ in E . We denote the configurations above C_1, C_2, C_3 and C_4 respectively, with respective stacks s_1, s_2, s_3, s_4 .

- Finally, B is a map assigning each head h in the graph a set $B(h)$ of *stack descriptors*, which are $(n+1)$ -tuples (h_n, \dots, h_1, h_c) of heads. In the following, we refer to h_k as the order- k component and h_c the collapse component. Roughly speaking, h_k describes at which head the new top_k -stack resulting from a pop_k operation (applied to a configuration with head h) may have been created, and h_c does likewise for a *collapse* operation. (We will use \perp in place of a head to indicate when pop_k or *collapse* is undefined.)

Consider $C_3 = \langle p_3, s_3 \rangle$ from the example above. This has control-state p_3 and top stack symbol c and so is associated with the head (p_3, c) . Thus $B((p_3, c))$ should contain the stack-descriptor $((p_1, b), (p_2, b), (p_1, b))$, which describes s_3 . The first (order-2) component is because $top_2(s_3)$ was created by a $push_2$ operation from a configuration with head (p_1, b) . The second (order-1) component is because the top symbol was created via an order-1 push from (p_2, b) . Finally, the order-2 link from the top of s_3 points to a stack occurring on top of a configuration at the head (p_1, b) , giving rise to the final (collapse) component describing the collapse link.

Tracking this information allows the summary algorithm to process the rule r_3 to obtain a description of C_4 from the description of C_3 . Since this rule performs a pop_1 , it can look at the order-1 component of the stack descriptor to see the head (p_2, b) , telling us that pop_1 results in b being on top of the stack. Since the rule r_3 moves into control-state p_4 , this tells us that the new head should be (p_4, b) . It also tells us that certain pieces of information in $B((p_2, b))$ are relevant to the description of $top_2(s_4)$ contained in $B((p_4, b))$. First remark that this situation only occurs for the pop_k and $collapse_k$ operations. To keep track of these correlations, we will introduce in Section 6.2.2 another component U of the graph.

More formally, Let us fix an ordinary n -CPDS with rules \mathcal{R} and initial configuration $c_0 := (p_0, [\dots[a_0]\dots])$. A *head* is an element $(p, a) \in \mathcal{P} \times \Sigma$ and should be viewed as

describing *stacks* u such that there is a *reachable* configuration of the form (p, u) where $top_1(u) = a$. Formally we define:

$$\llbracket (p, a) \rrbracket := \{u \in Stacks_n \mid top_1(u) = a \text{ and } (p, u) \in Post_C^*\}$$

A *stack descriptor* is an $(n+1)$ -tuple (h_n, \dots, h_1, h_c) where for each $1 \leq i \leq n$, each of h_i and h_c is either a head or \perp . We write $\mathbf{SDesc} := ((\mathcal{P} \times \Sigma) \cup \{\perp\})^{n+1}$ for the set of stack descriptors and it will also be useful to have $\mathbf{SDesc}_k := ((\mathcal{P} \times \Sigma) \cup \{\perp\})^{n-k}$ for the set of *order- k stack-descriptor prefixes*. Note that $\mathbf{SDesc}_n = \{()\}$ —i.e. consists only of the empty tuple. Assuming a map $B : (\mathcal{P} \times \Sigma) \rightarrow \mathbf{SDesc}$ a stack descriptor describes a set of stacks

$$\llbracket (h_n, \dots, h_1, h_c) \rrbracket := \left\{ u \in Stacks_n \mid \begin{array}{l} pop_k(u) \text{ is undefined if } h_k = \perp \text{ otherwise} \\ top_1(pop_k(u)) = b_k \text{ where } h_k = (_, b_k) \text{ and} \\ pop_k(u) \in \llbracket (h_n, \dots, h_{k+1}, h'_k, \dots, h'_1, h'_c) \rrbracket \\ \text{for some } (_, \dots, _, h'_k, \dots, h'_1, h'_c) \in B(h_k), \\ \text{for every } 1 \leq k \leq n, \\ \text{and } top_1(u) \text{ emits no link if } h_c = \perp \text{ otherwise} \\ top_1(collapse_k(u)) = b_c \text{ where } h_c = (_, b_c) \text{ and} \\ collapse_k(u) \in \llbracket (h_n, \dots, h_{k+1}, h'_k, \dots, h'_1, h'_c) \rrbracket \\ \text{for some } (_, \dots, _, h'_k, \dots, h'_1, h'_c) \in B(h_c), \\ \text{for every } 1 \leq k \leq n \end{array} \right\}$$

We now define an approximate reachability graph.

Definition 6.3. An *approximate reachability graph* for the CPDS \mathcal{C} is a triple (H, E, B) such that

- (i) $H \subseteq \mathcal{P} \times \Sigma$ is a set of heads such that $(p, u) \in Post_C^*$ implies that $(p, top_1(u)) \in H$,
- (ii) $E \subseteq H \times \mathcal{R} \times H$ is a set of triples such that if $(p, u) \in Post_C^*$ and
 - (a) $r = (p, top_1(u), o, p') \in \mathcal{R}$ for which $o(u)$ is defined, then it is the case that $((p, top_1(u)), r, (p', top_1(o(u)))) \in E$, and
 - (b) $r = (p, P) \in \mathcal{R}$ then $((p, top_1(u)), r, (p', top_1(u))) \in E$ for all $p' \in P$,
- (iii) B is a map $B : H \rightarrow \mathbf{SDesc}$ such that for every $h \in H$ we have $\llbracket h \rrbracket \subseteq \{\llbracket d \rrbracket \mid d \in B(h)\}$.

A non-trivial approximate reachability graph is computed using an algorithm that works *forwards* (while saturation works backwards), and which resembles a summary algorithm in the spirit of Sharir and Pnueli [29].

6.2.2. The Approximate Summary Algorithm. The construction of the approximate reachability graph is described in Algorithms 3, 4, 5 and 6. The main work is done in the function `ProcessHeadWithDescriptor`. In particular, this is where summary edges are added for the pop_k and $collapse_k$ operations.

The approximate summary algorithm computes an approximate reachability graph (H, E, B) ‘as accurately as possible based on an order-1 approximation’. In order to do this, the algorithm builds up an object (H, E, B, U) where the additional component U is a set of *approximate higher-order summary edges*. A summary edge describes how information contained in stack descriptors should be shared between heads. An *order- k summary edge* from a head h to a head h' is a triple of the form $(h, (h'_n, \dots, h'_{k+1}), h')$ where each h'_i is a head. That is, a triple in $H \times \mathbf{SDesc}_k \times H$.

Such a summary edge is added when processing either a pop_k or a $collapse_k$ operation on an order- k link. Intuitively such a summary edge means that if $(h_n, \dots, h_{k+1}, h_k, \dots, h_1, h_c) \in B(h)$, then we have $(h'_n, \dots, h'_{k+1}, h_k, \dots, h_1, h_c) \in B(h')$. When $n = k = 1$ (so that h_c is also unnecessary since there would be no links) note that $(h, (), h')$ behaves like a summary edge in a standard order-1 summary algorithm [29], which is complete at order-1.

To continue our example, the r_3 rule (which performs a pop_1 operation) from C_3 to C_4 means U should contain an order-1 summary edge $((p_2, b), ((p_1, b)), (p_4, b))$. Since pop_1 is an order-1 operation, we have $pop_2(s_3) = pop_2(s_4)$. Hence (p_1, b) (the order-2 component of the stack descriptor for s_3) should also be the first component of a stack descriptor for s_4 . However, since $top_1(s_4)$ was created at a configuration with head (p_2, b) , the order-1 and collapse components of such a stack descriptor for s_4 should be inherited from a stack descriptor in $B((p_2, b))$. In general if we go from a configuration (p, s) with head h to a configuration (p', s') with head h' by the pop_k operation or $collapse_k$ on an order- k link, we have that $pop_{k+1}(s) = pop_{k+1}(s')$ and hence we have a summary edge $(h, (h'_n, \dots, h'_{k+1}), h')$

The algorithm is presented as Algorithm 3.

Algorithm 3 The Approximate Summary Algorithm

Require: An n -CPDS with rules \mathcal{R} and heads $\mathcal{P} \times \Sigma$ and initial configuration $(p_0, [\dots [a_0] \dots])$

Ensure: The creation of a structure (H, E, B, U) where (H, E, B) is an approximate reachability graph and U is a set of approximate higher-order summary edges.

Set $H := \{(p_0, a_0)\}$ and set E, B and U to be empty

Call $\text{AddStackDescriptor}((p_0, a_0), (\perp, \dots, \perp, \perp))$

return Done, (H, E, B, U) will now be as required

Algorithm 4 $\text{AddStackDescriptor}(h, (h_n, \dots, h_1, h_c))$

Require: A head $h \in H$ and a stack descriptor (h_n, \dots, h_1, h_c)

Ensure: $(h_n, \dots, h_1, h_c) \in B(h)$ and that any further additions to $B(h')$ for each $h' \in H$ necessary to respect summary edges are made.

if $(h_n, \dots, h_1, h_c) \in B(h)$ **then**

return Done (Nothing to do)

Add (h_n, \dots, h_1, h_c) to $B(h)$

Call $\text{ProcessHeadWithDescriptor}(h, (h_n, \dots, h_1, h_c))$

for $h' \in H$ such that $(h, (h'_n, \dots, h'_{k+1}), h') \in U$ **do**

 Call $\text{AddStackDescriptor}(h', (h'_n, \dots, h'_{k+1}, h_k, \dots, h_1, h_c))$

return Done

Lemma 6.4. *Algorithm 3 terminates and the resulting structure (H, E, B, U) gives an approximate reachability graph (H, E, B) .*

Proof. For termination note that the respective procedures in Algorithms 4 and 6 will immediately return if the stack-descriptor (respectively summary) that they are called with is already contained in a particular set. If it does not belong to this set, then it is added. Since there are only finitely many possible arguments for these functions, they can thus only be called finitely many times without immediately returning. From this fact it is easy to see that the entire algorithm must always terminate.

Algorithm 5 ProcessHeadWithDescriptor($h, (h_n, \dots, h_1, h_c)$)

Require: A head $h := (p, a) \in H$ and a stack descriptor $(h_n, \dots, h_1, h_c) \in B(h)$
Ensure: All necessary modifications to the graph are made so that it is consistent with the presence of $(h_n, \dots, h_1, h_c) \in B(h)$. In particular this is the procedure that processes the CPDS rules from h (with respect to a stack described by h and the stack descriptor)
for o and p' such that $r = (p, a, o, p') \in \mathcal{R}$ **do**
 if o of form rew_b **then**
 Add (p', b) to H and $((p, a), r, (p', b))$ to E
 Call AddStackDescriptor($(p', b), (h_n, \dots, h_1, h_c)$)
 else if o of form $push_b^k$ **then**
 Add (p', b) to H and $((p, a), r, (p', b))$ to E
 Call AddStackDescriptor($(p', b), (h_n, \dots, h_2, (p, a), h_k)$)
 else if o of form $push_k$ **then**
 Add (p', a) to H and $((p, a), r, (p', a))$ to E
 Call AddStackDescriptor($(p', a), (h_n, \dots, h_{k+1}, (p, a), h_{k-1}, \dots, h_1, h_c)$)
 else if o of form pop_k with $h_k = (p_k, a_k)$ where $a_k \neq \perp$ **then**
 Add (p', a_k) to H and $((p, a), r, (p', a_k))$ to E
 Call AddSummary($(p_k, a_k), (h_n, \dots, h_{k+1}), (p', a_k)$)
 else if o of form $collapse_k$ with $h_c = (p_c, a_c)$ where $a_c \neq \perp$ **then**
 Add (p_c, a_c) to H and $((p, a), r, (p', a_c))$ to E
 Call AddSummary($(p_c, a_c), (h_n, \dots, h_{k+1}), (p', a_c)$)
for $r = p \rightarrow P \in \mathcal{R}$, and $p' \in P$ **do**
 Add (p', a) to H and $((p, a), r, (p', a))$ to E
 Call AddStackDescriptor($(p', a), (h_n, \dots, h_1, h_c)$)
return Done

Algorithm 6 AddSummary($h, (h'_n, \dots, h'_{k+1}), h'$)

Require: An approximate higher-order summary edge $(h, (h'_n, \dots, h'_{k+1}), h')$
Ensure: $(h, (h'_n, \dots, h'_{k+1}), h') \in U$ and that all necessary stack descriptors are added to the appropriate $B(h'')$ for $h'' \in H$ so that all summary edges (including the new one) are respected.
if $(h, (h'_n, \dots, h'_{k+1}), h') \in U$ **then**
 return Done (Nothing to do)
Add $(h, (h'_n, \dots, h'_{k+1}), h')$ to U
for $(h_n, \dots, h_{k+1}, h_k, \dots, h_1, h_c) \in B(h)$ **do**
 AddStackDescriptor($h', (h'_n, \dots, h'_{k+1}, h_k, \dots, h_1, h_c)$)
return Done

Now we show that (H, E, B) is an approximate reachability graph. Recursively define $Post_C^0 := \{c_0\}$ and

$$Post_C^{i+1} := Post_C^i \cup \{c \mid \exists c' \in Post_C^i \text{ s.t. } c' \rightarrow c \text{ or } c' \rightarrow C \text{ with } c \in C\}$$

That is $Post_C^i$ is the set of configurations that can be reached from the initial configuration in at most i steps. For a head $(p, a) \in \mathcal{P} \times \Sigma$, define

$$\llbracket (p, a) \rrbracket_i := \{(p, u) \mid (p, u) \in Post_C^i \text{ and } top_1(u) = a\}$$

We can now define an *i*-partial approximate reachability graph to be a version of an approximate reachability graph defined for ‘reachability up to depth *i*’.

Definition 6.5. An *i*-partial approximate reachability graph for the CPDS \mathcal{C} is a triple (H, E, B) such that

- (i) $H \subseteq \mathcal{P} \times \Sigma$ is a set of heads such that $(p, u) \in \text{Post}_{\mathcal{C}}^i$ implies that $(p, \text{top}_1(u)) \in H$,
- (ii) $E \subseteq H \times \mathcal{R} \times H$ is a set of triples such that if $i > 0$ and $(p, u) \in \text{Post}_{\mathcal{C}}^{i-1}$ and
 - (a) $r = (p, \text{top}_1(u), o, p') \in \mathcal{R}$ for which $o(u)$ is defined, then it is the case that $((p, \text{top}_1(u)), r, (p', \text{top}_1(o(u)))) \in E$,
 - (b) $r = (p, P) \in \mathcal{R}$ then $((p, \text{top}_1(u)), r, (p', \text{top}_1(u))) \in E$ for all $p' \in P$,
- (iii) B is a map $B : H \rightarrow \mathbf{SDesc}$ such that for every $h \in H$ we have $\llbracket h \rrbracket_i \subseteq \{\llbracket d \rrbracket \mid d \in B(h)\}$.

Observe that a structure (H, E, B) is an approximate reachability graph if and only if it is an *i*-partial approximate reachability graph for every $i \geq 0$.

Now observe that the algorithm monotonically grows the sets making up (H, E, B, U) (it only adds to the sets, it never removes from them). We may thus argue by induction to show that the (H, E, B) after termination is an *i*-partial approximate reachability graph for every $i \geq 0$ (and hence an approximate reachability graph). First note that the opening statements of Algorithm 3 (including the call to add $(\perp, \dots, \perp, \perp)$ as a stack descriptor to $B(p_0, a_0)$) guarantees that (H, E, B) is a 0-partial approximate reachability graph.

Now suppose (H, E, B) is an *i*-partial approximate reachability graph. We show it is also an $(i+1)$ -partial approximate reachability graph. Let $(p, u) \in \text{Post}_{\mathcal{C}}^i$ and let either

- (1) $r := (p, a, o, p') \in \mathcal{R}$ be such that $o(u)$ is defined and $\text{top}_1(u) = a$ so that $(p', o(u)) \in \text{Post}_{\mathcal{C}}^{i+1}$, or
- (2) $r := p \rightarrow P \in \mathcal{R}$ so that $(p', u) \in \text{Post}_{\mathcal{C}}^{i+1}$ for all $p' \in P$.

Let $a' := \text{top}_1(o(u))$. It suffices to show that (i) $h' := (p', a') \in H$, (ii) $e := ((p, a), r, (p', a')) \in E$ and that (iii) Some $d' := (h'_n, \dots, h'_1, h'_c) \in B(p', a')$ with $o(u) \in \llbracket d' \rrbracket$.

By the induction hypothesis (that the structure is an *i*-partial approximate reachability graph) we must have $h := (p, \text{top}_1(u)) \in H$ and $d = (h_n, \dots, h_1, h_c) \in B(h)$ such that $u \in \llbracket d \rrbracket$. Inspection of the algorithm shows that the addition of d to $B(h)$ is only possible if **AddStackDescriptor** (h, d) was called at some point during its execution. However, this also implies that **ProcessHeadWithDescriptor** (h, d) must have been called.

Note also that when o is a rewrite operation or r is an alternating rule we must have $\text{pop}_j(o(u)) = \text{pop}_j(u)$ and $\text{collapse}_j(o(u)) = \text{collapse}_j(u)$ for all j . When $o = \text{push}_k$ for $k \geq 2$ we must have $\text{top}_{j+1}(\text{pop}_j(o(u))) = \text{top}_{j+1}(\text{pop}_j(u))$ and $\text{top}_{j+1}(\text{collapse}_j(o(u))) = \text{top}_{j+1}(\text{collapse}_j(u))$ for all $j \neq k$ and $\text{pop}_k(o(u)) = u$. When $o = \text{push}_b^k$ we must have $\text{pop}_j(o(u)) = \text{pop}_j(u)$ for all $j \geq 2$, but $\text{pop}_1(o(u)) = u$ and $\text{collapse}_k(o(u)) = \text{pop}_k(u)$.

Thus if o is any operation other than pop_k or collapse_k it can be seen that the function **AddStackDescriptor** (h', d') must be called for a d' such that $u \in \llbracket d' \rrbracket$. Also, e is added to E . Since the algorithm never deletes elements from sets, this ensures that (H, E, B) must satisfy the constraints (i), (ii) and (iii) above.

Now consider the case when o is either pop_k or collapse_k . Suppose again that $\text{top}_1(o(u)) = a'$. Since $u \in \llbracket d \rrbracket$ we must have:

- For some control-state p^- we have: $h_k = (p^-, a')$ if $o = \text{pop}_k$ and $h_c = (p^-, a')$ if $o = \text{collapse}_k$ such that...
- ...there exists $(-, \dots, -, h'_k, \dots, h'_1, h'_c) \in B((p^-, a'))$ such that $o(u) \in \llbracket (h_n, \dots, h_{k+1}, h'_k, \dots, h'_1, h'_c) \rrbracket$.

Thus a suitable d' is $d' = (h_n, \dots, h_{k+1}, h'_k, \dots, h'_1, h'_c)$.

The call to **ProcessHeadWithDescriptor** (h, d) guarantees that (i) $h' = (p', a') \in H$ and (ii) $e := ((p, a), r, (p', a')) \in E$. It just remains to check that $d' \in B((p', a'))$.

Note that the above call must also ensure a call to

AddSummary $((p^-, a'), (h_n, \dots, h_{k+1}), (p', a'))$.

We are thus guaranteed the existence of a summary edge s : $((p^-, a'), (h_n, \dots, h_{k+1}), (p', a')) \in U$ (although it may have been added at an earlier point in the algorithm). There are two cases to consider:

- If the summary edge s was created *after* a stack-descriptor $(-, \dots, -, h'_k, \dots, h'_1, h'_c)$ was added to $B((p^-, a'))$, then the call to **AddSummary** creating s must add d' to $B((p', a'))$.
- If the summary edge s was created *before* a stack-descriptor $d^- = (-, \dots, -, h'_k, \dots, h'_1, h'_c)$ was added to $B((p^-, a'))$, then **ProcessHeadWithDescriptor** $((p^-, d^-))$ call creating this stack-descriptor must result in d' being added to $B((p', a'))$.

Either way, (iii) must also be satisfied. □

6.2.3. A Remark On Complexity. The approximate summary algorithm runs in time polynomial in the size of the CPDS (see below). Since the graph constructed must also be of polynomial size, it follows that the rules for the guarded CPDS \mathcal{C}' can also be extracted in polynomial time. Since the raw saturation algorithm is also PTIME when the number of control-states is fixed, it follows that the C-SHORE algorithm as a whole – including the forwards approximation and saturation – is fixed-parameter tractable.

We sketch here how to see that the approximate summary algorithm runs in polynomial time (when, as is standard, the order n is fixed). First note that an approximate reachability graph can contain at most $|Q| \cdot |\Sigma|$ heads and at most $|Q| \cdot |\Sigma| \cdot |\mathcal{R}| \cdot |Q| \cdot |\Sigma|$ edges (recalling that \mathcal{R} is the set of CPDS rules). Moreover the maximum size of the function B (when viewed as a relation defined by $\{(h, d) \in (Q \times \Sigma) \times (Q \times \Sigma)^{n+1} \mid d \in B(h)\}$) is $|Q| \cdot |\Sigma| \cdot (|Q| \cdot |\Sigma|)^{n+1}$. The maximum number of summary edges is $\sum_{i=2}^n |Q| \cdot |\Sigma| \cdot (|Q| \cdot |\Sigma|)^{n-i} \cdot |Q| \cdot |\Sigma|$. It follows that the size of the structure (H, E, B, U) constructed by algorithm is at most polynomial in the size of the original CPDS. Moreover, since the algorithm only *adds* to the structure and never removes elements previously added, it will perform at most polynomially many additions. Let Z be this polynomial bound on the size of the structure.

Moreover, recall that the procedures for adding summaries and heads/stack-descriptors are guarded. I.e. the procedure only processes the new object if it had not already been added; if it had already been added, the procedure will return after constant time.

So we consider the cases when the created object is new. For each *new* head/stack-descriptor pair, **ProcessHeadWithDescriptor** will check it against every rule and for each rule may attempt to create a new object. Disregarding the result of the calls to create *new* objects (with calls to create old objects returning in constant time), the run-time of this procedure will thus be bounded by $O(|\mathcal{R}|)$. Likewise each time a *new* stack descriptor is added, **AddStackDescriptor** will compare it against existing summary edges and so run in time $O(Z)$.

Similarly the run-time of a call to **AddSummary** on a new summary edge (disregarding run-times to calls from this procedure that create *new* objects) is $O(Z)$ since the new summary edge will, at worst, be compared against every possible stack-descriptor.

Thus creating a new object takes at most $O(Z \cdot |\mathcal{R}|)$ time and new objects are created only during the call to a procedure that itself is creating a new object. Thus the overall run-time is bounded by $O(Z \cdot Z \cdot |\mathcal{R}|)$ and so is polynomial.

6.3. Extracting the Guarded CPDA. Let $\mathcal{G} = (H, E, B)$ be an approximate reachability graph for \mathcal{C} . Let $\mathbf{Heads}(\mathcal{E})$ be the set of heads of error configurations, *i.e.* $\mathbf{Heads}(\mathcal{E}) := \{(p_{err}, a) \mid a \in \Sigma\}$. We do a simple backwards reachability computation on the finite graph \mathcal{G} to compute $\mathbf{BackRules}(\mathcal{G})$, defined to be the smallest set satisfying:

$$\begin{aligned} \mathbf{BackRules}(\mathcal{G}) = & \{e \in E \mid e = (h, r, h') \in E \text{ for some } h' \in \mathbf{Heads}(\mathcal{E})\} \cup \\ & \{e \in E \mid e = (h, r, h') \in E \text{ for some } (h', _, _) \in \mathbf{BackRules}(\mathcal{G})\} \end{aligned}$$

The CPDS rules occurring in the triples in $\mathbf{BackRules}(\mathcal{G})$ can be used to define a pruned CPDS that is safe if and only if the original also is. However, the approximate reachability graph provides enough information to construct a guarded CPDS whose guards are non-trivial. It is clear that the following set $\mathbf{BackRulesG}(\mathcal{G})$ of *guarded* rules can be computed:

$$\begin{aligned} & \{p \rightarrow P \mid (_, p \rightarrow P, _) \in \mathbf{BackRules}(\mathcal{G})\} \\ & \cup \\ & \left\{ (p, a, o', p') \mid \begin{array}{l} (_, (p, a, o, p'), _) \in \mathbf{BackRules}(\mathcal{G}) \text{ and} \\ o' = \begin{cases} o^S & \text{if } o \text{ is a pop or a collapse and } S \text{ is} \\ & \left\{ b \in \Sigma \mid ((p, a), r, (p', b)) \in E \right\} \\ & \text{with } r = (p, a, o, p') \\ o & \text{if } o \text{ is a rewrite or push} \end{cases} \end{array} \right\} \end{aligned}$$

These rules define a GCPDS on which C-SHORE finally performs saturation.

Lemma 6.6. *The GCPDS \mathcal{C}' defined using the rules $\mathbf{BackRulesG}(\mathcal{G})$ satisfies:*

$$Post_{\mathcal{C}}^* \cap Pre_{\mathcal{C}}^*(\mathcal{E}) \subseteq Pre_{\mathcal{C}'}^*(\mathcal{E}) \subseteq Pre_{\mathcal{C}}^*(\mathcal{E})$$

Proof. $Pre_{\mathcal{C}'}^*(\mathcal{E}) \subseteq Pre_{\mathcal{C}}^*(\mathcal{E})$ is trivial since $\mathbf{Triv}(\mathcal{C}')$ is a subset of the rules for \mathcal{C} .

Now suppose that $(p, u) \in Post_{\mathcal{C}}^* \cap Pre_{\mathcal{C}}^*(\mathcal{E})$. By (i) in the definition of approximate reachability graphs it must be the case that $(p, top_1(u)) \in H$ (since $(p, u) \in Post_{\mathcal{C}}^*$).

Since $(p, u) \in Pre_{\mathcal{C}}^*(\mathcal{E})$ we must also have $(p, u) \in Pre_{\mathcal{C}}^\alpha(\mathcal{E})$ where the proof of Lemma 5.4 gives the definition of $Pre_{\mathcal{C}}^\alpha(\mathcal{E})$. That is, (p, u) reaches \mathcal{E} in α steps. We induct over α .

When $\alpha = 0$ we have $(p, u) \in \mathcal{E}$ and the result is immediate. Otherwise, for $(\alpha+1)$, there are two cases. When $(p, u) \rightarrow C \subseteq Pre_{\mathcal{C}}^\alpha(\mathcal{E})$ via a rule $r = p \rightarrow P$ we have by induction $C \subseteq Pre_{\mathcal{C}'}^*(\mathcal{E})$ and by (i) and (ii) in the definition of approximate reachability graph, $h := (p, top_1(u)) \in H$ and $(h, r, h') \in E$ for every $h' := (p', top_1(u))$ with $p' \in P$. In the second case we have $(p, u) \rightarrow (p', o(u))$ by a rule $r := (p, a, o, p')$ and by induction $(p', o(u)) \in Pre_{\mathcal{C}'}^*(\mathcal{E})$. Furthermore, by (i) and (ii) in the definition of approximate reachability graphs we have $h := (p, top_1(u)) \in H$ and $(h, r, h') \in E$ where $h' := (p, top_1(o(u)))$.

Thus when r is alternating or o is neither a *pop* nor *collapse* operation $r' := r$ will itself occur as a rule of \mathcal{C}' . Otherwise $r' := (p, a, o^S, p')$ will be in \mathcal{C}' where $a \in S$. Thus applying r' to (p, u) witnesses $(p, u) \in Pre_{\mathcal{C}'}^*(\mathcal{E})$, as required. \square

7. EFFICIENT FIXED POINT COMPUTATION

We introduce an efficient method of computing the fixed point in Section 3, inspired by Schwoon *et al.*'s algorithm for alternating (order-1) pushdown systems [30]. Rather than checking all CPDS rules at each iteration, we fully process *all* consequences of each new transition at once. New transitions are kept in a set Δ_{new} (implemented as a stack), processed, then moved to a set Δ_{done} , which forms the transition relation of the final stack automaton. We assume w.l.o.g. that a character's link order is determined by the character. This is true for all CPDSs obtained from HORSSs.

In most cases, new transitions only depend on a single existing transition, hence processing the consequences of a new transition is straightforward. The key difficulty is the push rules and the alternating rules, which depend on *sets* of existing transitions. For example, given a rule $(p, a, push_k, p')$, processing $t = q_{p'} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_k, \dots, Q_n)$ 'once and once only' must somehow include adding a new transition whenever there is a set of transitions of the form $Q_k \xrightarrow[Q'_{col}]{a} (Q'_1, \dots, Q'_k)$ in A_i either now *or in the future*. When t is processed, we create a *trip-wire*, consisting of a *source* and a *target*. A *target* collects transitions from a given set of states (such as Q_k above), whilst a *source* describes how such a collection could be used to form a new transition according to a *push* saturation step.

Definition 7.1. An *order- k source* for $k \geq 1$ is a tuple (q_k, q_{k-1}, a, Q_k) in $\mathbb{Q}_k \times (\mathbb{Q}_{k-1} \cup \{\perp\}) \times \Sigma \times 2^{\mathbb{Q}_k}$. An *order- k target* is a tuple

$$\begin{cases} (Q_k, Q_k^C, Q_{lbl}, Q'_k) \in 2^{\mathbb{Q}_k} \times 2^{\mathbb{Q}_k} \times 2^{\mathbb{Q}_{k-1}} \times 2^{\mathbb{Q}_k} & \text{if } k \geq 2, \\ (Q_1, Q_1^C, a, Q_{col}, Q'_1) \in \bigcup_{k'=2}^n (2^{\mathbb{Q}_1} \times 2^{\mathbb{Q}_1} \times \Sigma \times 2^{\mathbb{Q}_{k'}} \times 2^{\mathbb{Q}_1}) & \text{if } k = 1. \end{cases}$$

The set Q_k^C is a *countdown* containing states in Q_k still awaiting a transition. We always have $Q_k^C \subseteq Q_k$ and $(Q_k \setminus Q_k^C) \xrightarrow{Q_{lbl}} Q'_k$. Likewise, an order-1 target $(Q_1, Q_1^C, a, Q_{col}, Q'_1)$ will satisfy $(Q_1 \setminus Q_1^C) \xrightarrow[Q_{col}]{a} Q'_1$. A target is *complete* if $Q_k^C = \emptyset$ or $Q_1^C = \emptyset$.

A *trip-wire* of order- k is an order- k source-target pair of the form $((-, -, -, Q_k), (Q_k, -, -, -))$ when $k \geq 2$ or $((-, -, a, Q_k), (Q_k, -, a, -, -))$ when $k = 1$. When the target in a trip-wire is *complete*, the action specified by its source is triggered, which we now sketch.

An order- k *source* for $k \geq 2$ describes how an order- $(k-1)$ source should be created from a complete target, propagating the computation to the level below, and an order-1 source describes how a new long-form transition should be created from a complete target. That is, when we have $(q_k, -, a, Q_k)$ (we hide the second component for simplicity of description) and $(Q_k, \emptyset, Q_{lbl}, Q'_k)$ this means we have found a set of transitions witnessing $Q_k \xrightarrow{Q_{lbl}} Q'_k$ and should now look for transitions from Q_{lbl} . Hence the algorithm creates a new source and target for the order- $(k-1)$ state-set Q_{lbl} . When this process reaches order-1, a new transition is created. This results in the construction of the t' from a *push* saturation step.

Let the function **extract_short_forms** obtain from a long-form transition its (unique) corresponding set of (*short-form*) transitions.

Algorithm 7 gives the main loop and introduces the global sets Δ_{done} and Δ_{new} , and two arrays $\mathcal{U}_{src}[k]$ and $\mathcal{U}_{targ}[k]$ containing sources and targets for each order. The algorithm processes pop_n and $collapse_n$ rules like the naive algorithm and creates trip wires for the alternating transitions. Algorithm 8 gives the main steps processing a new transition. In

Algorithm 7 Computing $Pre_C^*(A_0)$

Let $\Delta_{done} = \emptyset$, $\Delta_{new} = \bigcup_{n \geq k \geq 1} \Delta_k$, $\mathcal{U}_{src}[k] = \emptyset$, $\mathcal{U}_{targ}[k] = \{(\emptyset, \emptyset, \emptyset, \emptyset)\}$ for each $n \geq k > 1$ and $\mathcal{U}_{targ}[1] = \{(\emptyset, \emptyset, a, \emptyset, \emptyset) \mid a \in \Sigma\}$.

for $r := (p, a, pop_n, p') \in \mathcal{R}$ **do**

$add_to_worklist\left(q_p \xrightarrow[\emptyset]{a} (\emptyset, \dots, \emptyset, \{q_{p'}\}), r\right)$

for $r := (p, a, collapse_n, p') \in \mathcal{R}$ **do**

$add_to_worklist\left(q_p \xrightarrow[\{q_{p'}\}]{a} (\emptyset, \dots, \emptyset), r\right)$

for $r := p \rightarrow P \in \mathcal{R}$ and $a \in \Sigma$ **do**

$create_trip_wire(q_p, \perp, Q, r)$ where $Q = \{q_{p'} \mid p' \in P\}$

while $\exists t \in \Delta_{new}$ **do**

$update_rules(t)$

$update_trip_wires(t)$

Move t from Δ_{new} to Δ_{done}

Algorithm 8 $update_rules(t)$

Require: A transition t to be processed against Δ_{done}

if t is an order- k transition for $2 \leq k \leq n$ of the form $q_{p'Q_n \dots Q_{k+1}} \rightarrow Q_k$ **then**

for $p \in \mathcal{P}$ and $a \in \Sigma$ such that $r := (p, a, pop_{k-1}, p') \in \mathcal{R}$ **do**

$add_to_worklist\left(q_p \xrightarrow[Q_{col}]{a} (\emptyset, \dots, \emptyset, \{q_{p'Q_n \dots Q_k}\}, Q_k, \dots, Q_n), r\right)$

for $p \in \mathcal{P}$ and $a \in \Sigma$ such that $r := (p, a, collapse_{k-1}, p') \in \mathcal{R}$ **do**

$add_to_worklist\left(q_p \xrightarrow[\{q_{p'Q_n \dots Q_k}\}]{a} (\emptyset, \dots, \emptyset, Q_k, \dots, Q_n), r\right)$

for $p \in \mathcal{P}$ and $a \in \Sigma$ such that $r := (p, a, push_k, p') \in \mathcal{R}$ **do**

$create_trip_wire(q_{p, Q_n, \dots, Q_{k+1}}, q_{p', Q_n, \dots, Q_{k+1}, Q_k}, a, Q_k, (r, t))$

else if t is an order-1 transition of the form $q_{p'Q_n \dots Q_2} \xrightarrow[Q_{col}]{b} Q_1$ **then**

for $p \in \mathcal{P}$ and $a \in \Sigma$ such that $r := (p, a, rew_b, p') \in \mathcal{R}$ **do**

$add_to_worklist\left(q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n), (r, t)\right)$

for $p \in \mathcal{P}$ and $a \in \Sigma$ such that $r := (p, a, push_b^k, p') \in \mathcal{R}$ **do**

$create_trip_wire(q_{p, Q_n, \dots, Q_{k+1}, Q_k \cup Q_{col}, Q_{k-1}, \dots, Q_2}, \perp, a, Q_1, (r, t))$

Algorithm 9 $update_trip_wires\left(t = q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)\right)$

for $t_k = q_k \rightarrow Q_k \in \mathbf{extract_short_forms}(t)$ **do**

for $targ \in \mathcal{U}_{targ}[k]$ with $targ = (-, Q_{k'}^C, -, -)$ or $(-, Q_{k'}^C, a, -, -)$ and $q_k \in Q_{k'}^C$ **do**

$proc_targ_against_tran(targ, t_k)$

Algorithm 10 $\text{create_trip_wire}(q_k, q_{k-1}, a, Q_k, jus)$

```

if  $(q_k, q_{k-1}, a, Q_k) \notin \mathcal{U}_{src}[k]$  then
  Add  $src := (q_k, q_{k-1}, a, Q_k)$  to  $\mathcal{U}_{src}[k]$ , set  $J(src) := jus$ 
  Let  $targ := (Q_k, Q_k, \emptyset, \emptyset)$  if  $k > 1$  or  $(Q_k, Q_k, a, \emptyset, \emptyset)$  if  $k = 1$ 
  if  $targ \in \mathcal{U}_{targ}[k]$  then
    for each complete target  $targ$  matching  $src$  do
       $\text{proc\_source\_complete\_targ}(src, targ)$ 
  else
     $\text{add\_target}(targ, k)$ ; set  $J(targ) := \emptyset$ 

```

Algorithm 11 $\text{proc_targ_against_tran}(targ, t)$

```

Suppose  $\begin{cases} t = q_k \xrightarrow{\quad} Q_k'' \text{ and } targ = (Q_k, Q_k^C, Q_{lbl}, Q_k') & \text{if } k \geq 2 \\ t = q_1 \xrightarrow[Q_{col}]{a} Q_1'' \text{ and } targ = (Q_1, Q_1^C, a, Q_{lbl}, Q_1') & \text{if } k = 1 \end{cases}$ 

Let  $targ' := \begin{cases} (Q_k, Q_k^C \setminus \{q_k\}, Q_{lbl} \cup \{q_k Q_k''\}, Q_k' \cup Q_k'') & \text{if } k \geq 2 \\ (Q_1, Q_1^C \setminus \{q_1\}, a, Q_{lbl} \cup Q_{col}, Q_1' \cup Q_1'') & \text{if } k = 1 \end{cases}$ 

if  $q_k \in Q_k^C$  and  $targ' \notin \mathcal{U}_{targ}[k]$  then
   $\text{add\_target}(targ', k)$ ; if  $k = 1$ , set  $J(targ') := J(targ) \cup \{t\}$ 
  if  $Q_k^C \setminus \{q_k\} = \emptyset$  then
    for each source  $src \in \mathcal{U}_{src}[k]$  of form  $(-, -, -, Q_k)$  do
       $\text{proc\_source\_complete\_targ}(src, targ')$ 

```

Algorithm 12 $\text{proc_source_complete_targ}(src, comp_targ)$

Require: An order- k source of the form $src = (q_k, q_{k-1}, a, Q_k)$ and an order- k complete target of the form $comp_targ = (Q_k, \emptyset, Q_{lbl}, Q_k')$ when $k \geq 2$ and $(Q_1, \emptyset, a, Q_{lbl}, Q_1')$ when $k = 1$

```

if  $k \geq 2$  then
   $\text{create\_trip\_wire}(q_k Q_k', \perp, a, Q_{lbl} \cup S, J(src))$  where  $S := \{q_{k-1} \mid q_{k-1} \neq \perp\}$ 
else if  $k = 1$  then
  Suppose  $q_1 = q_p, Q_n, \dots, Q_2$  and  $J(comp\_targ) = T$ 
   $\text{add\_to\_worklist}(q_p \xrightarrow[Q_{lbl}]{a} (Q_1', Q_2, \dots, Q_n), jus)$  where  $jus = \begin{cases} (r, t, T) & J(src) = (r, t) \\ (r, T) & J(src) = r \end{cases}$ 

```

Algorithm 13 $\text{add_to_worklist}(t, jus)$

Require: A long form transition t and justification jus .

```

for  $u \in \text{extract\_short\_forms}(t)$  such that  $u \notin \Delta_{done} \cup \Delta_{new}$  do
  Add  $u$  to  $\Delta_{new}$  and set  $J(u) := (jus, |\Delta_{new} \cup \Delta_{done}|)$  if  $u$  is order-1

```

Algorithm 14 $\text{add_target}(targ, k)$

```

if  $targ \notin \mathcal{U}_{targ}[k]$  then
  Add  $targ$  to  $\mathcal{U}_{targ}[k]$ 
  for  $t' \in \Delta_{done}$  do
     $\text{proc\_targ\_against\_tran}(targ, t')$ 

```

most cases a new transition is created, however, for *push* rules we create a trip-wire. We describe some of the algorithms informally below.

In `create_trip_wire` we create a trip-wire with a new target $(Q_k, Q_k, \emptyset, \emptyset)$. This is added using an `add_target` procedure which also checks Δ_{done} to create further targets. E.g., a new target $(Q_k, Q_k^C, Q_{lbl}, Q'_k)$ combines with an existing $q_k \xrightarrow{q_{k-1}} Q''_k$ to create a new target $(Q, Q_k^C \setminus \{q_k\}, Q_{lbl} \cup \{q_{k-1}\}, Q'_k \cup Q''_k)$. (This step corrects a bug of Schwoon *et al.*) Similarly `update_trip_wires` updates existing targets by new transitions. In all cases, when a source and matching complete target are created, we perform the propagations as above.

Proposition 7.2. *Given a CPDS \mathcal{C} and stack automaton A_0 , let A be the result of Algorithm 7. We have $\mathcal{L}(A) = Pre_{\mathcal{C}}^*(A_0)$. \square*

7.1. Correctness. We prove Proposition 7.2. I.e., the fast algorithm is correct. The proof is in two parts in the following sub-sections. In particular in Lemma 7.10 and Lemma 7.5.

In the sequel, we fix the following notation. Let $(A_i)_{i \geq 0}$ be the sequence of automata constructed by the naive fixed point algorithm. Then, let $(\Delta_{done}^j)_{j \geq 0}$ be the sequence of sets of transitions such that Δ_{done}^j is Δ_{done} after j iterations of the main loop of Algorithm 7. Similarly, define $\mathcal{U}_{src}^j[k]$ and $\mathcal{U}_{targ}^j[k]$.

7.1.1. Soundness. We prove that the algorithm is sound. First, we show two preliminary lemmas about the data-structures maintained by the algorithm.

Lemma 7.3. *For all $j \geq 0$ and $n \geq k > 1$, if $(Q_k, Q_k \setminus Q_k^T, Q_{k-1}^T, Q_k^{T'}) \in \mathcal{U}_{targ}^j[k]$, then we have $T \subseteq \Delta_{done}^j$ that witnesses $Q_k^T \xrightarrow{Q_{k-1}^T} Q_k^{T'}$.*

Proof. We proceed by induction over j and the order in which targets are created. In the base case we only have $(\emptyset, \emptyset, \emptyset, \emptyset) \in \mathcal{U}_{targ}^0[k]$. Setting $T = \emptyset$ witnesses $\emptyset \xrightarrow{\emptyset} \emptyset$.

In the inductive case, consider the location of the call to `add_target`: `create_trip_wire` or `proc_targ_against_tran`. When in `create_trip_wire`, we have a target of the form $(Q_k, Q_k, \emptyset, \emptyset)$, hence $Q_k^T = \emptyset$ and we trivially have $T = \emptyset \subseteq \Delta_{done}^j$ witnessing $\emptyset \xrightarrow{\emptyset} \emptyset$.

Otherwise the call is from `proc_targ_against_tran` against a transition $t = q_k \rightarrow Q''_k$ and a target $targ = (Q_k, Q_k \setminus Q_k^T, Q_{k-1}^T, Q_k^{T'})$ already in \mathcal{U}_{targ} . Hence, by induction, we

know that there is some $T \subseteq \Delta_{done}^j[k]$ witnessing $Q_k^T \xrightarrow{Q_{k-1}^T} Q_k^{T'}$. The transition t is either already in Δ_{done}^j or will be moved there at the end of the j th iteration. Combining t with

T we have $T \cup \{t\} \subseteq \Delta_{done}^j$ witnessing $Q_k^T \cup \{q_k\} \xrightarrow{Q_{k-1}^T \cup \{q_k Q''_k\}} Q_k^{T'} \cup Q''_k$. Since the new target added is $(Q_k, Q_k \setminus (Q_k^T \cup \{q_k\}), Q_{k-1}^T \cup \{q_k Q''_k\}, Q_k^{T'} \cup Q''_k)$ we are done. \square

Lemma 7.4. *For all $j \geq 0$, if $(Q_1, Q_1 \setminus Q_1^T, a, Q_{col}^T, Q_1^{T'}) \in \mathcal{U}_{targ}^j[1]$, then we have $T \subseteq \Delta_{done}^j$ that witnesses $Q_1^T \xrightarrow[Q_{col}^T]{a} Q_1^{T'}$.*

Proof. The proof is essentially the same as the order- k case above. We proceed by induction over j and the order in which targets are created. In the base case we only have $(\emptyset, \emptyset, a, \emptyset, \emptyset) \in \mathcal{U}_{targ}^0[1]$. Setting $T = \emptyset$ witnesses $\emptyset \xrightarrow[\emptyset]{a} \emptyset$.

In the inductive case, consider the location of the call to `add_target`: `create_trip_wire` or `proc_targ_against_tran`. When in `create_trip_wire`, we have a target of the form $(Q_1, Q_1, a, \emptyset, \emptyset)$, hence $Q_1^T = \emptyset$ and we trivially have $T = \emptyset \subseteq \Delta_{done}^j$ witnessing $\emptyset \xrightarrow[\emptyset]{a} \emptyset$.

Otherwise the call is from `proc_targ_against_tran` against a transition $t = q_1 \xrightarrow[Q_{col}^T]{a} Q_1''$ and a target $targ = (Q_1, Q_1 \setminus Q_1^T, Q_{col}^T, Q_1^{T'})$ already in \mathcal{U}_{targ} . Hence, by induction, we know that there is some $T \subseteq \Delta_{done}^j[1]$ witnessing $Q_1^T \xrightarrow[Q_{col}^T]{a} Q_1^{T'}$. The transition t is either already in Δ_{done}^j or will be moved there at the end of the j th iteration. Combining t with T we have $T \cup \{t\} \subseteq \Delta_{done}^j$ witnessing $Q_1^T \cup \{q_1\} \xrightarrow[Q_{col}^T \cup Q_{col}]{a} Q_1^{T'} \cup Q_1''$. Since the new target is $(Q_1, Q_1 \setminus (Q_1^T \cup \{q_1\}), a, Q_{col} \cup Q_{col}^T, Q_1^{T'} \cup Q_1'')$ we are done. \square

We are now ready to prove the algorithm is sound.

Lemma 7.5. *Given a CPDS \mathcal{C} and stack automaton A_0 , let A be the result of Algorithm 7. We have $\mathcal{L}(A) \subseteq Pre_{\mathcal{C}}^*(A_0)$.*

Proof. We proceed by induction over j and show every transition appearing in Δ_{done}^j appears in A_i for some i . This implies the lemma.

When $j = 0$ the property is immediate, since the only transitions added are already in A_0 , or added to A_1 during the first processing of the `popn` and `collapsen` rules.

In the inductive step, we consider some t first appearing in Δ_{new}^j (and thus, eventually in $\Delta_{done}^{j'}$ for some j'). There are several cases depending on how t was added to Δ_{new} (i.e. from where `add_to_worklist` was called). We consider the simple cases first. In all the following cases, t was added during `update_rules` against a transition t' appearing in Δ_{new}^{j-1} .

- If $t' = q_{p'Q_n \dots Q_{k+1}} \longrightarrow Q_k$ and t was added as part of

$$t_1 = q_p \xrightarrow[Q_{col}]{a} (\emptyset, \dots, \emptyset, \{q_{p'Q_n \dots Q_k}\}, Q_k, \dots, Q_n)$$

during the processing of t' against a `popk-1` rule. By induction t' appears in A_i for some i , and hence t_1 (which includes t) is present in A_{i+1} .

- If $t' = q_{p'Q_n \dots Q_{k+1}} \longrightarrow Q_k$ and t was added as part of

$$t_1 = q_p \xrightarrow[\{q_{p'Q_n \dots Q_k}\}]{a} (\emptyset, \dots, \emptyset, Q_k, \dots, Q_n)$$

during the processing of t' against a `collapsek-1` rule. By induction t' appears in A_i for some i , and hence t_1 (which includes t) is present in A_{i+1} .

- If $t' = q_{p'Q_n \dots Q_2} \xrightarrow[Q_{col}]{b} Q_1$ and t was added as part of $t_1 = q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ during the processing of t' against a rew_b rule. By induction t' appears in A_i for some i , and hence t_1 (which includes t) is present in A_{i+1} .

In the final case, `add_to_worklist` is called during `proc_source_complete_targ`. There are two cases depending on the provenance of the source. In the first case, the source was added by a call to `create_trip_wire` from `update_rules` while processing a $push_b^k$ rule against $t' = q_{p'Q_n \dots Q_2} \xrightarrow[Q_{col}]{b} Q_1$. Therefore, t was added as part of

$$q_p \xrightarrow[Q_{col}]{a} (Q'_1, Q_2, \dots, Q_{k-1}, Q_k \cup Q_{col}, Q_{k+1}, Q_n)$$

from a source $(q_1, \perp, a, Q_1) \in \mathcal{U}_{src}^j[1]$ with $q_1 = q_{p, Q_n, \dots, Q_{k+1}, Q_k \cup Q_{col}, Q_{k-1}, \dots, Q_2}$. By induction, from t' we know that $q_{p'} \xrightarrow[Q_{col}]{b} (Q_1, \dots, Q_n)$ appears in A_i for some i . Now, consider the target $(Q_1, \emptyset, a, Q'_{col}, Q'_1) \in \mathcal{U}_{targ}^j[1]$ that was combined with the source to add the new transition. By Lemma 7.4 we have $Q_1 \xrightarrow[Q'_{col}]{a} Q'_1$ in Δ_{done}^j and hence (since all transitions in Δ_{done}^j passed through Δ_{new}) by induction we have $Q_1 \xrightarrow[Q'_{col}]{a} Q'_1$ in $A_{i'}$ for some i' . Hence, in $A_{\max(i, i') + 1}$ we have t as required.

In the second case we have a source $(q_1, \perp, a, Q_1^s) \in \mathcal{U}_{src}^j[1]$ and a complete target of the form $(Q_1^s, \emptyset, a, Q_{col}^t, Q_1^{t'}) \in \mathcal{U}_{targ}^j[1]$ and the source derived from a call to `create_trip_wire` in `proc_source_complete_targ`. Note, by Lemma 7.4 we have $Q_1^s \xrightarrow[Q_{col}^t]{a} Q_1^{t'}$ in Δ_{done}^j . The call to `create_trip_wire` implies we have a source $(q_2, q'_1, a, Q_2^s) \in \mathcal{U}_{src}^j[2]$ and complete target of the form $(Q_2^s, \emptyset, Q_1^t, Q_2^{t'}) \in \mathcal{U}_{targ}^j[2]$, with $Q_1^s = Q_1^t \cup S_1$ where $S_1 = \{q'_1 \mid q'_1 \neq \perp\}$ and $q_1 = q_2 Q_2^{t'}$. The proof will now iterate from $k = 2$ upwards until a source is discovered that was added during a call to `create_trip_wire` from `update_rules` while processing some $push_{k'}$ rule or alternating rule. Note that sources not added by $push_b^k$ rules can only be added in this way and, for all $k < k'$, the second component of the source (q'_{k-1}) will be \perp .

Hence, inductively, we have a source $src = (q_k, q'_{k-1}, a, Q_k^s) \in \mathcal{U}_{src}^j[k]$ and complete target $(Q_k^s, \emptyset, Q_{k-1}^t, Q_k^{t'}) \in \mathcal{U}_{targ}^j[k]$ with $Q_{k-1}^s = Q_{k-1}^t \cup S_{k-1}$ where $S_{k-1} = \{q'_{k-1} \mid q'_{k-1} \neq \perp\}$ and $q_{k-1} = q_k Q_k^{t'}$. Furthermore, by Lemma 7.3 we have $Q_k^s \xrightarrow[Q_{col}^t]{a} Q_k^{t'}$ in Δ_{done}^j .

In the first case, suppose src was added due to a call to `create_trip_wire` in a call to `proc_source_complete_targ`. The call to `create_trip_wire` implies we have a source of the form $(q_{k+1}, q'_k, a, Q_{k+1}^s) \in \mathcal{U}_{src}^j[k+1]$ and complete target $(Q_{k+1}^s, \emptyset, Q_k^t, Q_{k+1}^{t'}) \in \mathcal{U}_{targ}^j[k+1]$, with $Q_k^s = Q_k^t \cup S_k$ where $S_k = \{q'_k \mid q'_k \neq \perp\}$ and $q_k = q_{k+1} Q_{k+1}^{t'}$.

For the final cases, first suppose that src was added due to a call to `create_trip_wire` in `update_rules` from a $push_k$ rule. Then we were processing a new transition $q_{p'Q_n \dots Q_{k+1}} \longrightarrow Q_k$, and we have $q_k = q_{p, Q_n, \dots, Q_{k+1}}$ and $q'_{k-1} = q_{p', Q_n, \dots, Q_k}$ and $Q_k^s = Q_k$. From the induction and since $q'_{k'} = \perp$ for all $k' < k$, we have $Q_{k-1}^t \cup \{q'_{k-1}\} \xrightarrow[Q_{col}^t]{a} (Q_1^{t'}, \dots, Q_{k-1}^{t'})$ in Δ_{done}^j .

which can be split into $Q_{k-1}^t \xrightarrow[Q'_{col}]{a} (Q'_1, \dots, Q'_{k-1})$ and $q'_{k-1} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_{k-1})$. Thus, because $q'_{k-1} = q_{p', Q_n, \dots, Q_k}$ and $Q_k^s = Q_k$ and letting $Q'_k = Q_k^{t'}$, we have

$$q_{p'} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n) \quad \text{and} \quad Q_k \xrightarrow[Q'_{col}]{a} (Q'_1, \dots, Q'_k)$$

in Δ_{done}^j and thus by induction in A_i for some i . Since we have

$$q_1 = q_{kQ'_k, Q_{k-1} \cup Q'_{k-1}, \dots, Q_1 \cup Q'_1} = q_{p, Q_n, \dots, Q_{k+1}, Q'_k, Q_{k-1} \cup Q'_{k-1}, \dots, Q_2 \cup Q'_2}$$

we added t as part of a transition

$$q_p \xrightarrow[Q_{col} \cup Q'_{col}]{a} (Q_1 \cup Q'_1, \dots, Q_{k-1} \cup Q'_{k-1}, Q'_k, Q_{k+1}, \dots, Q_n)$$

which is the transition added by the naive saturation algorithm from the $push_k$ rule and the transitions in A_i . Hence, we satisfy the lemma.

Otherwise, src was added due to a call to `create_trip_wire` during initialisation from an alternating rule $p \rightarrow P$. Then we have $k = n$, $q_k = q_p$, $q'_{k-1} = \perp$, and $Q_k^s = \{q_{p'} \mid p' \in P\}$. From the induction we have $Q_{k-1}^t \xrightarrow[Q_{col}^t]{a} (Q'_1, \dots, Q'_{k-1})$ in Δ_{done}^j . Thus we have $Q_k^s \xrightarrow[Q_{col}^t]{a} (Q'_1, \dots, Q'_n)$ in Δ_{done}^j and thus by induction in A_i for some i . We added t as part of $q_p \xrightarrow[Q_{col}^t]{a} (Q'_1, \dots, Q'_n)$ which is the transition added by the naive saturation algorithm given the alternating rule and transitions in A_i . \square

7.1.2. Completeness. We prove that the algorithm is complete. For this we need some preliminary lemmas stating properties of the data-structures maintained by the algorithm.

Lemma 7.6. *For all $k \geq 2$ and $j \geq 0$, all $T \subseteq \Delta_{done}^j$ witnessing $Q_k^T \xrightarrow[Q'_{k-1}]{Q_k^T} Q'_k$, and all $(q_k, q_{k-1}, a, Q_k) \in \mathcal{U}_{src}^j[k]$ such that $Q_k^T \subseteq Q_k$, we have $(Q_k, Q_k \setminus Q_k^T, Q_k^T, Q'_k)$ in $\mathcal{U}_{targ}^j[k]$.*

Proof. Let j_1 be the iteration of Algorithm 7 where (q_k, q_{k-1}, a, Q_k) was first added to $\mathcal{U}_{src}^{j_1}[k]$. We perform an induction over j_1 . The base case is trivial. In the inductive case, the only position where a source may be added is in the `create_trip_wire` procedure. After adding the source, we re-establish the induction hypothesis. There are two cases.

Let $targ = (Q_k, Q_k, \emptyset, \emptyset)$. If $targ$ is already in $\mathcal{U}_{targ}^{j_1}$ then we observe that a target of the form (Q_k, Q_k, \dots) is only created in `create_trip_wire` (targets are also created in `proc_targ_against_tran`, but these targets are obtained by removing a state from the second component of an existing target, hence the two first components cannot be equal). This implies there exists a source $(_, _, _, Q_k) \in \mathcal{U}_{src}^{j'}[k]$ for some $j' < j_1$. This gives the result by induction since T and the desired target depend only on the final component of the source.

If $targ$ is not in $\mathcal{U}_{targ}^{j_1}[k]$, then we add it. Next, split $T = T_1 \cup T_2$ such that T_1 contains all $t \in T$ appearing in $\Delta_{done}^{j_1-1}$. The balance is contained in T_2 . The algorithm proceeds to call `proc_targ_against_tran` on $targ$ and all $t \in \Delta_{done}^{j_1}$. In particular, this includes all $t \in T_1$.

We aim to prove that, after the execution of this loop, we have $(Q_k, Q_k \setminus Q_k^1, Q_k^1, Q_k^{1'}) \in \mathcal{U}_{targ}^{j_1}[k]$ when T_1 witnesses $Q_k^1 \xrightarrow[Q_{k-1}^1]{Q_k^1} Q_k^{1'}$.

Let t_1, \dots, t_ℓ be a linearisation of T_1 in the order they appear in iterations over Δ_{done} (we assume a fixed order here for convenience, though the proof can generalise if the order changes between iterations). Additionally, let $T_z = \{t_1, \dots, t_z\}$ witness $Q_k^{t_z} \xrightarrow{Q_{k-1}^{t_z}} Q_k^{t'_z}$. We show after T_z has been processed, we have $(Q_k, Q_k \setminus Q_k^{t_z}, Q_{k-1}^{t_z}, Q_k^{t'_z}) \in \mathcal{U}_{targ}^{j_1}[k]$. This gives us the property once $z = \ell$. In the base case $z = 0$ and we are done. Otherwise, we know $targ_z = (Q_k, Q_k \setminus Q_k^{t_z}, Q_{k-1}^{t_z}, Q_k^{t'_z}) \in \mathcal{U}_{targ}^{j_1}[k]$ and prove the case for $(z+1)$. Consider the call to `add_target` that added $targ_z$. Now take the iteration against Δ_{done} that processes t_{z+1} . This results in the addition of $(Q_k, Q_k \setminus Q_k^{t_{z+1}}, Q_{k-1}^{t_{z+1}}, Q_k^{t'_{z+1}})$ as required.

Hence, we have $(Q_k, Q_k \setminus Q_k^1, Q_{k-1}^1, Q_k^{1'}) \in \mathcal{U}_{targ}^{j_1}[k]$. Now, let t_1, \dots, t_ℓ be a linearisation of T_2 in the order they are added to Δ_{done} . Additionally, we write $Q_k^{t_z} \xrightarrow{Q_{k-1}^{t_z}} Q_k^{t'_z}$ for the state-sets and transitions witnessed by $T_1 \cup \{t_1, \dots, t_z\}$.

We show after t_z has been added to Δ_{done} on the j' th iteration, we have that $(Q_k, Q_k \setminus Q_k^{t_z}, Q_{k-1}^{t_z}, Q_k^{t'_z}) \in \mathcal{U}_{targ}^{j'}[k]$ for some j' . In the base case $z = 0$ and we are done by the argument above. Otherwise, we know $targ_z = (Q_k, Q_k \setminus Q_k^{t_z}, Q_{k-1}^{t_z}, Q_k^{t'_z}) \in \mathcal{U}_{targ}^{j'}[k]$ and prove the case for $(z+1)$. Consider the call to `update_trip_wires` with t_{z+1} . This results in the addition of the target $(Q_k, Q_k \setminus Q_k^{t_{z+1}}, Q_{k-1}^{t_{z+1}}, Q_k^{t'_{z+1}})$ via the call to `proc_targ_against_tran`. When $z = \ell$, we have the lemma as required. \square

Lemma 7.7. *For all $j \geq 0$, all $T \subseteq \Delta_{done}^j$ witnessing $Q_1^T \xrightarrow[Q_{col}^T]{a} Q_1'$, and all $(q_1, \perp, a, Q_1) \in \mathcal{U}_{src}^j[1]$ such that $Q_k^T \subseteq Q_1$, we have $(Q_1, Q_1 \setminus Q_1^T, a, Q_{col}^T, Q_1')$ in $\mathcal{U}_{targ}^j[1]$.*

Proof. The proof is essentially the same as the proof when $k \geq 2$. Let j_1 be the iteration of Algorithm 7 where (q_1, q_{k-1}, a, Q_1) was first added to $\mathcal{U}_{src}^{j_1}[1]$. We perform an induction over j_1 . In the base case the lemma is trivially true. In the inductive case, the only position where a source may be added is in the `create_trip_wire` procedure. After adding the source, the induction hypothesis needs to be re-established. There are two cases.

Let $targ = (Q_1, Q_1, a, \emptyset, \emptyset)$. If $targ$ is already in $\mathcal{U}_{targ}^{j_1}$ then we observe that a target of the form (Q_1, Q_1, \dots) is only created in `create_trip_wire`. This implies the existence of a source $(_, _, _, Q_1) \in \mathcal{U}_{src}^{j'}[1]$ for some $j' < j_1$. This implies the result by induction since neither T nor the desired target depend any but the final component of the source.

If $targ$ is not in $\mathcal{U}_{targ}^{j_1}[1]$, then we add it. Next, split $T = T_1 \cup T_2$ such that T_1 contains all $t \in T$ appearing in $\Delta_{done}^{j_1-1}$. The balance is contained in T_2 . The algorithm proceeds to call `proc_targ_against_tran` on $targ$ and all $t \in \Delta_{done}^{j_1}$. In particular, this includes all $t \in T_1$.

We aim to prove that, after the execution of this loop, we have $(Q_1, Q_1 \setminus Q_1^1, a, Q_{col}^1, Q_1^{1'}) \in \mathcal{U}_{targ}^{j_1}[1]$ when T_1 witnesses $Q_1^1 \xrightarrow[Q_{col}^1]{a} Q_1^{1'}$.

Let t_1, \dots, t_ℓ be a linearisation of T_1 in the order they appear in iterations over Δ_{done} . Additionally, let $T_z = \{t_1, \dots, t_z\}$ witness $Q_1^{t_z} \xrightarrow[Q_{col}^{t_z}]{a} Q_1^{t'_z}$. We show after T_z has been processed, we have $(Q_1, Q_1 \setminus Q_1^{t_z}, a, Q_{col}^{t_z}, Q_1^{t'_z}) \in \mathcal{U}_{targ}^{j_1}[1]$. This gives us the property once $z = \ell$. In the base case $z = 0$ and we are done. Otherwise, we know $targ_z = (Q_1, Q_1 \setminus Q_1^{t_z}, a, Q_{col}^{t_z}, Q_1^{t'_z}) \in \mathcal{U}_{targ}^{j_1}[1]$ and prove the case for $(z+1)$. Consider the call to `add_target`

that added $targ_z$. Now take the iteration against Δ_{done} that processes t_{z+1} . This results in the addition of $(Q_1, Q_1 \setminus Q_1^{t_{z+1}}, a, Q_{col}^{t_{z+1}}, Q_1^{t_{z+1}})$ as required.

Hence, we have $(Q_1, Q_1 \setminus Q_1^1, a, Q_{col}^1, Q_1^1) \in \mathcal{U}_{targ}^{j_1}[1]$. Now, let t_1, \dots, t_ℓ be a linearisation of T_2 in the order they are added to Δ_{done} . Additionally, we write $Q_1^{t_z} \xrightarrow[Q_{col}^{t_z}]{a} Q_1^{t'_z}$ for the state-sets and transitions witnessed by $T_1 \cup \{t_1, \dots, t_z\}$.

We show after t_z is added to Δ_{done} on the j' 'th iteration, we have that $(Q_1, Q_1 \setminus Q_1^{t_z}, a, Q_{col}^{t_z}, Q_1^{t'_z}) \in \mathcal{U}_{targ}^{j'}[1]$ for some j' . In the base case $z = 0$ and we are done by the argument above. Otherwise, we know $targ_z = (Q_1, Q_1 \setminus Q_1^{t_z}, a, Q_{col}^{t_z}, Q_1^{t'_z}) \in \mathcal{U}_{targ}^{j'}[1]$ and prove the case for $(z + 1)$. Consider the call to `update_trip_wires` with t_{z+1} . This results in the addition of the target $(Q_1, Q_1 \setminus Q_1^{t_{z+1}}, a, Q_{col}^{t_{z+1}}, Q_1^{t'_{z+1}})$ via the call to `proc_targ_against_tran`. When $z = \ell$, we have the lemma as required. \square

Lemma 7.8. *For all $k > 1$ and $j \geq 0$, if we have $(q_k, q_{k-1}, a, Q_k) \in \mathcal{U}_{src}^j[k]$ and also $(Q_k, \emptyset, Q_{k-1}, Q'_k) \in \mathcal{U}_{targ}^j[k]$, then it is the case that there exists $j' \geq 0$ such that we have $(q_k Q'_k, \perp, a, Q_{k-1} \cup S) \in \mathcal{U}_{src}^{j'}[k-1]$ where $S = \{q_{k-1} \mid q_{k-1} \neq \perp\}$.*

Proof. Let j_1 be the smallest such that $(q_k, q_{k-1}, a, Q_k) \in \mathcal{U}_{src}^{j_1}[k]$ and j_2 be the smallest such that $(Q_k, \emptyset, Q_{k-1}, Q'_k) \in \mathcal{U}_{targ}^{j_2}[k]$.

In the case $j_1 \leq j_2$, we consider the j_2 th iteration of Algorithm 7 at the moment where the target is added to $\mathcal{U}_{targ}^{j_2}[k]$. This has to be a result of the call to `add_target` during Algorithm 11. The only other place `add_target` may be called is during Algorithm 10; however, this implies the target is of the form $(Q_k, Q_k, \emptyset, \emptyset)$ and hence, for the target to be complete, it must be $(\emptyset, \emptyset, \emptyset, \emptyset)$ and hence $j_2 = 0$, and since $j_1 > 0$ (since there are initially no sources) we have a contradiction. Hence, the target is added during Algorithm 11 and the procedure goes on to call `proc_source_complete_targ` against each matching source in $\mathcal{U}_{src}^{j_2}[k]$, including (q_k, q_{k-1}, a, Q_k) . This results in the addition of $(q_k Q'_k, \perp, a, Q_{k-1} \cup S)$ to $\mathcal{U}_{src}^{j_2}[k-1]$, if it is not there already, satisfying the lemma.

In the case $j_1 > j_2$, we consider the j_1 th iteration of Algorithm 7 at the moment where the source is added. This is necessarily in the `create_trip_wire` procedure. Since $(Q_k, \emptyset, Q_{k-1}, Q'_k) \in \mathcal{U}_{targ}^{j_1}[k]$ and since this target must have been obtained from a target of the form $(Q_k, Q_k, \emptyset, \emptyset)$, we know $(Q_k, Q_k, \emptyset, \emptyset) \in \mathcal{U}_{targ}^{j_1}[k]$ and thus the procedure calls `proc_source_complete_targ` for each complete target including $(Q_k, \emptyset, Q_{k-1}, Q'_k)$. Thus we add $(q_k Q'_k, \perp, a, Q_{k-1} \cup S)$ to $\mathcal{U}_{src}^{j_2}[k-1]$, if it is not there already, satisfying the lemma. \square

Lemma 7.9. *For all $j \geq 0$, if $(q_1, \perp, a, Q_1) \in \mathcal{U}_{src}^j[1]$ and $(Q_1, \emptyset, Q_{col}, Q'_1) \in \mathcal{U}_{targ}^j[1]$, if $q_1 = q_{p, Q_n, \dots, Q_2}$, then for each t in*

$$\text{extract_short_forms}\left(q_p \xrightarrow[Q_{col}]{a} (Q'_1, Q_2, \dots, Q_n)\right)$$

there exists some $j' \geq 0$ such that $t \in \Delta_{done}^{j'}$.

Proof. As before, the proof of this order-1 case is very similar to the order- k proof.

Let j_1 be the smallest such that $(q_1, \perp, a, Q_1) \in \mathcal{U}_{src}^{j_1}[1]$ and j_2 be the smallest such that $(Q_1, \emptyset, a, Q_{col}, Q'_1) \in \mathcal{U}_{targ}^{j_2}[1]$.

In the case $j_1 \leq j_2$, we consider the j_2 th iteration of Algorithm 7 at the moment where the target is added to $\mathcal{U}_{\text{targ}}^{j_2}[1]$. This has to be a result of the call to `add_target` during Algorithm 11. The only other place `add_target` may be called is during Algorithm 10; however, this implies the target is of the form $(Q_1, Q_1, \emptyset, \emptyset)$ and hence, for the target to be complete, it must be $(\emptyset, \emptyset, \emptyset, \emptyset)$ and hence $j_2 = 0$, and since $j_1 > 0$ (since there are initially no sources) we have a contradiction. Hence, the target is added during Algorithm 11 and the procedure goes on to call `proc_source_complete_targ` against each matching source in $\mathcal{U}_{\text{src}}^{j_2}[1]$, including (q_1, \perp, a, Q_1) . This results in the addition of $q_p \xrightarrow[Q_{\text{col}}]{a} (Q'_1, Q_2, \dots, Q_n)$ satisfying the lemma.

In the case $j_1 > j_2$, we consider the j_1 th iteration of Algorithm 7 at the moment where the source is added. This is necessarily in the `create_trip_wire` procedure. Since $(Q_1, \emptyset, a, Q_{\text{col}}, Q'_1) \in \mathcal{U}_{\text{targ}}^{j_1}[1]$ and since this target must have been obtained from a target of the form $(Q_1, Q_1, a, \emptyset, \emptyset)$, we know $(Q_1, Q_1, a, \emptyset, \emptyset) \in \mathcal{U}_{\text{targ}}^{j_1}[1]$ and thus the procedure calls `proc_source_complete_targ` against each complete target including $(Q_1, \emptyset, a, Q_{\text{col}}, Q'_1)$. This results in the addition of the source $q_p \xrightarrow[Q_{\text{col}}]{a} (Q'_1, Q_2, \dots, Q_n)$ satisfying the lemma. \square

We are now ready to prove completeness.

Lemma 7.10. *Given a CPDSC and stack automaton A_0 , let A be the result of Algorithm 7. We have $\mathcal{L}(A) \supseteq \text{Pre}_C^*(A_0)$.*

Proof. We know (from the correctness of saturation (Theorem 5.1)) that the fixed point of $(A_i)_{i \geq 0}$ is an automaton recognising $\text{Pre}_C^*(A_0)$. We prove, by induction, that for each transition t appearing in A_i for some i , there exists some j such that t appears in Δ_{done}^j .

We first prove the only if direction. In the base case we have all transitions in A_0 in Δ_{new} at the beginning of Algorithm 7. Since the main loop continues until Δ_{new} has been completely transferred to Δ_{done} , the result follows.

Now, let t be an order- k transition appearing for the first time in A_i ($i > 0$). We case split on the pushdown operation that led to the introduction of the transition. Let $r = (p, a, o, p')$ be the rule that led to the new transition. We first consider simple cases.

- When $o = \text{pop}_k$, then there was $q_{p'} \xrightarrow{q_k} (Q_{k+1}, \dots, Q_n)$ in A_i and we added to A_{i+1}

$$q_p \xrightarrow[\emptyset]{a} (\emptyset, \dots, \emptyset, \{q_k\}, Q_{k+1}, \dots, Q_n)$$

of which t is a transition. By induction we have j such that $t' = q_{p'}Q_n \dots Q_{k+2} \xrightarrow{q_k} Q_{k+1}$ appears in Δ_{done}^j . Consider the j th iteration of Algorithm 7 when `update_rules` is called on t' . The `popk'` loop immediately adds

$$q_p \xrightarrow[\emptyset]{a} (\emptyset, \dots, \emptyset, \{q_k\}, Q_{k+1}, \dots, Q_n)$$

and hence t to Δ_{new} , giving us some $j' > j$ such that t appears in $\Delta_{\text{done}}^{j'}$.

- When $o = \text{collapse}_k$, when $k = n$, we added t as part of $q_p \xrightarrow[\{q_{p'}\}]{a} (\emptyset, \dots, \emptyset)$. In this case

we also added t to Δ_{new} as part of the initialisation Algorithm 7. Otherwise, $n > k$ and from a transition $q_{p'} \xrightarrow{q_k} (Q_{k+1}, \dots, Q_n)$ we added t as part of

$$q_p \xrightarrow[\{q_k\}]{a} (\emptyset, \dots, \emptyset, Q_{k+1}, \dots, Q_n)$$

By induction we have j such that $t' = q_{p'Q_n \dots Q_{k+2}} \xrightarrow{q_k} Q_{k+1}$ appears in Δ_{done}^j . Consider the j th iteration of Algorithm 7 when `update_rules` is called on t' . The `collapsek` loop immediately adds

$$q_p \xrightarrow[\{q_k\}]{a} (\emptyset, \dots, \emptyset, Q_{k+1}, \dots, Q_n)$$

and hence t to Δ_{new} , giving us some $j' > j$ such that t appears in $\Delta_{done}^{j'}$.

- when $o = rew_b$ then from a transition $q_{p'} \xrightarrow[Q_{col}]{b} (Q_1, \dots, Q_n)$ we added t as part of a transition $q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$. By induction, we know that $t' = q_{p', Q_n, \dots, Q_2} \xrightarrow[Q_{col}]{b} Q_1$ appears in Δ_{done}^j for some j . Consider the j th iteration of the main loop of Algorithm 7. During this iteration t' is passed to `update_rules`, and the loop handling rules containing `rewb` adds $q_p \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_n)$ to Δ_{new} . Since this transition contains t , there must be some j' such that t appears in $\Delta_{done}^{j'}$.

We now consider the push rules, which require more intricate reasoning.

- when $o = push_k$, we had $q_{p'} \xrightarrow[Q_{col}]{a} (Q_1, \dots, Q_k, \dots, Q_n)$ and T of the form $Q_k \xrightarrow[Q'_{col}]{a} (Q'_1, \dots, Q'_k)$ in A_i , and we added the transition

$$q_p \xrightarrow[Q_{col} \cup Q'_{col}]{a} (Q_1 \cup Q'_1, \dots, Q_{k-1} \cup Q'_{k-1}, Q'_k, Q_{k+1}, \dots, Q_n)$$

which has t . By induction, there is some j where $t_1 = q_{p'Q_n, \dots, Q_{k+1}} \rightarrow Q_k$ first appears in Δ_{done}^j . Also by induction, for each $t' \in T$, there is some j' such that t' first appears in $\Delta_{done}^{j'}$. We divide T into $T_1 \cup \dots \cup T_k$, where $T_{k'}$ contains all order- k' transitions in T .

Consider the j th iteration where t_1 is added to Δ_{done} . During the call to `update_rules` we call `create_trip_wire` in the loop handling push rules with $q_k = q_{p, Q_n, \dots, Q_{k+1}}$ and $q_{k-1} = q_{p', Q_n, \dots, Q_{k+1}}$, $a = a$ and $Q_k = Q_k$.

The call ensures $(q_k, q_{k-1}, a, Q_k) \in \mathcal{U}_{src}^j[k]$. Now take j' such that all $T_k \in \Delta_{done}^{j'}$. We know T_k witnesses $Q_k \xrightarrow[Q'_{col}]{Q_{k-1}} Q'_k$. By Lemma 7.6 we know that we have $(Q_k, \emptyset, Q_{k-1}, Q'_k) \in \mathcal{U}_{targ}^{j'}[k]$, and then additionally by Lemma 7.8 that we have $(q_{kQ'_k}, \perp, a, Q_{k-1} \cup \{q_{k-1}\}) \in \mathcal{U}_{src}^{j''}[k-1]$ for some j'' .

We iterate the above argument from $k' = k-1$ down to $k' = 1$. Begin with j' such that $(q_{k'}, \perp, a, Q_{k'}^{lbl}) \in \mathcal{U}_{src}^{j'}[k']$ and all $T_{k'} \in \Delta_{done}^{j'}$. We know $T_{k'}$ witnesses $Q_{k'}^{lbl} \xrightarrow[Q'_{k'}]{Q_{k'-1}^{lbl}} Q_{k'} \cup Q'_{k'}$. By Lemma 7.6 we know it to be the case that $(Q_{k'}^{lbl}, \emptyset, Q_{k'-1}^{lbl}, Q_{k'} \cup Q'_{k'}) \in \mathcal{U}_{targ}^{j'}[k]$, and then by Lemma 7.8 that we have $(q_{k'Q_{k'} \cup Q'_{k'}}, \perp, a, Q_{k'-1}^{lbl}) \in \mathcal{U}_{src}^{j''}[k-1]$ for some j'' .

When $k' = 1$, we have some j' with $(q_1, \perp, a, Q_1^{lbl}) \in \mathcal{U}_{src}^{j'}[1]$ and $T_1 \in \Delta_{done}^{j'}$. Note,

$$q_1 = q_{p, Q_n, \dots, Q_{k+1}, Q'_k, Q_{k-1} \cup Q'_{k-1}, \dots, Q_2 \cup Q'_2}.$$

We know $T_{k'}$ witnesses $Q_1^{lbl} \xrightarrow[Q_{col} \cup Q'_{col}]{a} Q_1 \cup Q'_1$. By Lemma 7.7 we also know that $(Q_1^{lbl}, \emptyset, a, Q_{col} \cup Q'_{col}, Q_1 \cup Q'_1) \in \mathcal{U}_{targ}^{j'}[k]$, and then by Lemma 7.9 we have j'' such that

we have all t' in

$$q_p \xrightarrow[Q_{col} \cup Q'_{col}]{a} (Q_1 \cup Q'_1, \dots, Q_{k-1} \cup Q'_{k-1}, Q'_k, Q_{k+1}, \dots, Q_n)$$

in $\Delta_{done}^{j''}$. This, in particular, includes t .

- when $o = push_b^k$ we had transitions $q_{p'} \xrightarrow[Q_{col}]{b} (Q_1, \dots, Q_n)$ and $T = Q_1 \xrightarrow[Q'_{col}]{a} Q'_1$ in A_i with $Q_{col} \subseteq Q_k$ and added the transitions (which include t)

$$q_p \xrightarrow[Q'_{col}]{a} (Q'_1, Q_2, \dots, Q_{k-1}, Q_k \cup Q_{col}, Q_{k+1}, \dots, Q_n) .$$

By induction, there is some j where $t_1 = q_{p'Q_n, \dots, Q_2} \xrightarrow[Q_{col}]{b} Q_1$ is first in Δ_{done}^j . In addition, for each $t' \in T$, there is some j' such that t' first appears in $\Delta_{done}^{j'}$.

Consider the j th iteration where t_1 is added to Δ_{done} . During the call to `update_rules` we call `create_trip_wire` in the loop handling push rules with

$$q_1 = q_k = q_{p, Q_n, \dots, Q_{k+1}, Q_k \cup Q_{col}, Q_{k-1}, \dots, Q_2}$$

$q_{k-1} = \perp$, $a = b$ and $Q_k = Q_1$.

The call ensures $(q_1, \perp, b, Q_1) \in \mathcal{U}_{src}^j[1]$. Now take j' such that all $T_1 \in \Delta_{done}^{j'}$. We know T witnesses $Q_1 \xrightarrow[Q_{col}]{a} Q'_1$. By Lemma 7.7 we know that $(Q_1, \emptyset, a, Q_{col}, Q'_1) \in \mathcal{U}_{targ}^{j'}[1]$, and then by Lemma 7.9 we have j'' such that we have all t' in

$$q_p \xrightarrow[Q'_{col}]{a} (Q_1 \cup Q'_1, Q_2, \dots, Q_{k-1}, Q_k \cup Q_{col}, Q_{k+1}, \dots, Q_n)$$

in $\Delta_{done}^{j''}$. This, in particular, includes t .

Finally, we consider the alternating rules. Take a rule $p \rightarrow P$. We had T of the form $Q \xrightarrow[Q'_{col}]{a} (Q'_1, \dots, Q'_n)$ in A_i , where $Q = \{q_{p'} \mid p' \in P\}$, and we added the transition $q_p \xrightarrow[Q'_{col}]{a} (Q'_1, \dots, Q'_n)$ which contains t . By induction, for each $t' \in T$, there is some j' such that t' first appears in $\Delta_{done}^{j'}$. We divide T into $T_1 \cup \dots \cup T_n$, where T_k contains all order- k transitions in T .

During initialisation we call `create_trip_wire` with $q_k = q_p$, $q_{k-1} = \perp$, $a = a$ and $Q_k = Q$. The call ensures $(q_k, q_{k-1}, a, Q_k) \in \mathcal{U}_{src}^j[k]$. Now take j' such that all $T_k \in \Delta_{done}^{j'}$. We know T_k witnesses $Q_k \xrightarrow[Q'_{col}]{a} Q'_k$. By Lemma 7.6 we know that we have $(Q_k, \emptyset, Q_{k-1}, Q'_k) \in \mathcal{U}_{targ}^{j'}[k]$, and then additionally by Lemma 7.8 that we have $(q_{kQ'_k}, \perp, a, Q_{k-1}) \in \mathcal{U}_{src}^{j''}[k-1]$ for some j'' . As above, we iterate the above argument down to $k' = 1$ until we have some j'' such that we have all t' in $q_p \xrightarrow[Q'_{col}]{a} (Q'_1, \dots, Q'_n)$ in $\Delta_{done}^{j''}$. This, in particular, includes t .

This completes the proof. \square

8. EXPERIMENTAL RESULTS

We compared C-SHORE with the state-of-the-art verification tools for higher-order recursion schemes (HORS) available on its release: TRecS [19], GTRecS2 [20] (the successor of [17]), and TravMC [21]. In an extension to our original publication [4], we have re-run these experiments to also compare with the verification tools released after C-SHORE: Preface [25], and HorSat2 [14]. As a further extension, we have tested the efficacy of the improved fixed point computation in Section 7 by implementing a naive fixed point computation where, during each iteration, each rule is tested against the current automaton to search for new transtions.

Benchmarks are from the TRecS and TravMC benchmark suites, plus several larger examples provided by Kobayashi. The majority of the TravMC benchmarks were translated into HORS from an extended formalism, HORS with Case statements (HORSC), using a script by Kobayashi. For fairness, all tools in our experiments took a pure HORS as input. However, the authors of TravMC report that TravMC performs faster on the original HORSC examples than on their HORS translations.

In all cases, the benchmarks consist of a HORS (generating a computation tree) and a property automaton. In the case of C-SHORE, the property automaton is a regular automaton describing branches of the generated tree that are considered errors. Thus, following the intuition in Section 2, we can construct a reachability query over a CPDS, where the reachability of a control state p_{err} indicates an erroneous branch (see [8] for more details). All other tools check co-reachability properties of HORS and thus the property automaton describes only valid branches of the computation tree. In all cases, it was straightforward to translate between the co-reachability and reachability properties.

The experiments were run on a Dell Latitude e6320 laptop with 4Gb of RAM and four 2.7GHz Intel i7-2620M cores. We ran C-SHORE on OpenJDK 8.0 using the argument “-Xmx” to limit RAM usage to 2.5Gb. As advised by the TravMC developers, we ran TravMC and Preface on the Mono JIT compiler (version 4.6.1) with no command line arguments. Finally TRecS (version 1.34), GTRecS2 (version 3.17), and HorSat2 were compiled with the OCaml version 4.02.3 compilers. On negative examples, GTRecS2 was run with its `-neg` argument. We used the “ulimit” command to limit memory usage to 2.5Gb and set a CPU timeout of 600 seconds (per benchmark). The given runtimes were reported by the respective tools and are the means of three separate runs on each example. Note, C-SHORE was run until the automaton was completely saturated.

Table 1 shows trials where at least one tool took over 1s. This is to save space and because virtual machine “warm-up” and HORS to CPDS conversion can skew the results on small benchmarks. Examples violating their property are marked “(bug)”. The order (Ord) and size (Sz) of the schemes were reported by TRecS. We show reported times in seconds for TRecS (T), GTRecS2 (G), TravMC (TMC), Preface (P), HorSat2 (H), and C-SHORE (C) as well as C-SHORE implementing a naive fixed point computation for the saturation (N). A dash “—” means analysis failed. In the next column we mark when C-SHORE was the fastest (✓) and slowest (✗) amongst its previous competitors (not including Preface or HorSat2). For C-SHORE, we then report the times for HORS to CPDS translation (Ctran), CPDS analysis (Ccpds), and building the approximation graph (Capprox). Capprox is part of Ccpds, and the full time (C) is the sum of Ctran and Ccpds.

Of 35 benchmarks, C-SHORE outperformed its previous competitors on 7 examples. In 9 cases, C-SHORE was the slowest, but in only 2 of those cases did C-SHORE require

Benchmark file	Ord	Sz	T	TMC	G	N	C	P	H	✓/✗	Ctran	Ccpds	Capprox
example3-1 (bug)	1	8	0.000	0.111	—	0.060	0.059	0.293	0.003		0.027	0.032	0.016
file	1	8	0.000	0.032	—	0.051	0.053	0.286	0.003		0.026	0.027	0.022
fileocamlc	4	111	0.027	0.047	0.042	—	0.222	0.295	0.010	✗	0.045	0.177	0.130
lock2	4	45	0.036	0.050	0.261	—	0.235	0.331	0.010		0.034	0.201	0.101
order5	5	52	0.013	0.042	—	37.152	0.250	0.315	0.010		0.037	0.213	0.090
order5-2	5	40	0.044	0.073	—	—	0.163	0.317	0.007		0.034	0.129	0.070
order5-variant	5	55	0.043	0.042	1.094	—	0.242	0.322	0.010		0.038	0.204	0.077
filepath	2	5956	215.401	—	—	0.205	0.212	0.503	0.040	✓	0.075	0.136	0.130
filter-nonzero (bug)	5	484	0.013	0.141	0.284	—	1.783	0.554	0.026	✗	0.064	1.719	1.450
filter-nonzero-1	5	890	0.281	96.163	—	—	5.018	1.827	0.100		0.093	4.925	4.244
map-head-filter (bug)	3	370	0.012	0.123	0.076	—	0.298	0.393	0.013	✗	0.055	0.243	0.093
map-head-filter-1	3	880	0.238	0.698	—	0.242	0.229	0.366	0.016	✓	0.071	0.158	0.151
map-plusone	5	302	0.034	0.088	0.224	—	0.827	0.398	0.013	✗	0.063	0.765	0.605
map-plusone-1	5	459	0.057	0.388	—	—	1.443	0.478	0.037		0.078	1.365	1.132
map-plusone-2	5	704	1.423	6.450	—	—	2.750	0.588	0.081		0.086	2.664	2.235
safe-head	3	354	0.048	0.046	0.040	—	0.246	0.364	0.012	✗	0.047	0.199	0.066
safe-init	3	680	0.081	0.147	0.263	—	0.486	0.416	0.016	✗	0.071	0.415	0.103
safe-tail	3	468	0.061	0.051	0.052	—	0.306	0.391	0.013	✗	0.058	0.248	0.093
g41	4	31	—	0.046	0.067	—	—	0.321	0.006	✗	0.027	—	0.116
cfa-life2	14	7648	—	—	—	—	—	0.857	0.173		0.431	—	—
cfa-matrix-1	8	2944	17.358	—	—	16.905	17.311	0.412	0.056	✓	0.225	17.086	17.081
cfa-psdes	7	1819	17.850	—	—	1.331	1.452	0.363	0.033	✓	0.143	1.309	1.301
dna	2	411	0.069	0.173	0.063	21.220	6.867	11.553	0.038	✗	0.120	6.746	6.303
exp4-5	4	55	—	—	0.306	—	—	0.389	0.010		0.032	—	2.410
fibstring	4	29	—	33.340	0.066	—	—	0.294	0.004		0.031	—	0.132
fold_fun_list	7	1346	0.618	—	—	1.262	1.284	0.327	0.020		0.109	1.175	1.169
fold_right	5	1310	32.123	—	—	1.248	1.335	0.331	0.021	✓	0.106	1.229	1.222
jwig-cal_main	2	7627	0.127	0.053	—	4.662	5.137	0.530	0.137		5.087	0.050	0.044
l	3	35	—	7.523	0.020	0.131	0.129	0.297	0.006		0.030	0.100	0.092
search-e-church (bug)	6	837	0.023	0.218	—	—	5.708	3.623	0.038		0.102	5.606	1.790
specialize_cps_coerce1-c	3	2731	—	—	—	0.463	0.503	0.433	0.206	✓	0.184	0.320	0.313
tak (bug)	8	451	—	2.002	—	—	50.032	3.276	0.090		0.084	49.948	42.078
xhtmlf-div-2 (bug)	2	3003	0.333	—	13.401	3.497	3.651	1.438	1.597		3.360	0.291	0.269
xhtmlf-m-church	2	3027	0.336	—	5.342	3.542	3.441	0.754	1.153		3.194	0.247	0.240
zip	4	2952	22.606	—	—	—	2.567	0.728	0.060	✓	0.157	2.409	1.612

Table 1: Comparison of model-checking tools.

more than 1 second. In general, both Preface and HorSat2 outperform all previous tools. It is worth noting that HorSat2, which appears to perform the best, is an adaptation of our saturation algorithm to recursion schemes [6].

Notably, C-SHORE does not perform well on `g41` and `exp4-5`. These belong to a class of benchmarks that stress higher-order model-checkers and indicate that our tool currently does not always scale well. However, C-SHORE seems to show a more promising capacity to scale on larger HORS produced by tools such as MoCHi [18], which are particularly pertinent in that they are generated by an actual software verification tool. We also note that C-SHORE timed out on the fewest examples of the previous tools despite not always terminating in the fastest time.

Finally, without the forwards analysis described in Section 6, all shown examples except `filepath` timed out. In addition, the naive version of the saturation algorithm performed significantly worse than the improved fixed point computation presented in Section 7.

9. CONCLUSION

We have given a full account of the C-SHORE tool. This includes the development of a *saturation* algorithm for CPDS that we first introduced in ICALP 2016 [3]. This is a backwards reachability algorithm. To produce a viable implementation we optimised this algorithm using two main approaches. The first is a preliminary forwards analysis which allows the input CPDS to be pruned and guarded, leading to faster analysis times. The second is an efficient fixed point computation. This implementation was first published in ICFP 2013 [4].

We have extended these results here by providing a generalisation of the implemented algorithms to alternating CPDS. Furthermore, we have implemented a naive version of the fixed point iteration required by saturation. Since this naive implementation is significantly out-performed by our efficient algorithm, we provide justification for the development in Section 7.

C-SHORE remains the only implementation of higher-order model checking using CPDS. This provides a completely novel approach which was competitive with its contemporary tools. Since its release, two new tools, Preface and HorSat (and HorSat2), were developed. These new tools are currently the fastest model-checkers for HORS.

Thanks. Robin Neatherway, Steven Ramsay, and Naoki Kobayashi for help with benchmarking, Łukasz Kaiser and Royal Holloway for web-hosting, and Stefan Schwoon. This work was supported by Fond. Sci. Math. Paris, AMIS [ANR 2010 JCJC 0203 01 AMIS], FREC [ANR 2010 BLAN 0202 02 FREC], VAPF (Région IdF), and the Engineering and Physical Sciences Research Council[EP/K009907/1].

REFERENCES

- [1] T. Ball and S. K. Rajamani. The SLAM project: Debugging system software via static analysis. In *POPL*, pages 1–3, Portland, Oregon, January 16–18, 2002.
- [2] A. Bouajjani and A. Meyer. Symbolic Reachability Analysis of Higher-Order Context-Free Processes. In *Proc. 24rd Conf. on Found. of Software Technology and Theoretical Computer Science (FSTTCS'04)*, volume 3328 of *Lecture Notes in Computer Science*, Madras, India, December 2004. Springer Pub.
- [3] C. H. Broadbent, A. Carayol, M. Hague, and O. Serre. A saturation method for collapsible pushdown systems. In *ICALP*, pages 165–176, 2012.
- [4] C. H. Broadbent, A. Carayol, M. Hague, and O. Serre. C-shore: a collapsible approach to higher-order verification. In *ICFP*, pages 13–24, 2013.

- [5] C. H. Broadbent, A. Carayol, C.-H. Luke Ong, and O. Serre. Recursion schemes and logical reflection. In *LICS*, pages 120–129, 2010.
- [6] C. H. Broadbent and N. Kobayashi. Saturation-based model checking of higher-order recursion schemes. In *CSL*, pages 129–148, 2013.
- [7] C-SHORE. <http://cshore.cs.rhul.ac.uk/>.
- [8] A. Carayol and O. Serre. Collapsible pushdown automata and labeled recursion schemes: Equivalence, safety and effective selection. In *LICS*, pages 165–174, 2012.
- [9] Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981.
- [10] M. Hague, A. S. Murawski, C.-H. Luke Ong, and O. Serre. Collapsible pushdown automata and recursion schemes. In *LICS*, pages 452–461, 2008.
- [11] M. Hague and C.-H. L. Ong. Symbolic backwards-reachability analysis for higher-order pushdown systems. *Logical Methods in Computer Science*, 4(4), 2008.
- [12] M. Hague and C.-H. L. Ong. Analysing mu-calculus properties of pushdown systems. In *SPIN*, pages 187–192, 2010.
- [13] N. D. Jones and S. S. Muchnick. Even simple programs are hard to analyze. *J. ACM*, 24:338–350, April 1977.
- [14] N. Kobayashi. HorSat2: A model checker for HORS based on SATuration. A tool available at <http://www-kb.is.s.u-tokyo.ac.jp/~koba/horsat2/>.
- [15] N. Kobayashi. Types and higher-order recursion schemes for verification of higher-order programs. In *POPL*, pages 416–428, 2009.
- [16] N. Kobayashi. Higher-order model checking: From theory to practice. In *LICS*, pages 219–224, 2011.
- [17] N. Kobayashi. A practical linear time algorithm for trivial automata model checking of higher-order recursion schemes. In *FOSSACS*, pages 260–274, 2011.
- [18] N. Kobayashi, R. Sato, and H. Unno. Predicate abstraction and cegar for higher-order model checking. In *PLDI*, pages 222–233, 2011.
- [19] Naoki Kobayashi. Model-checking higher-order functions. In *PPDP*, pages 25–36, 2009.
- [20] Naoki Kobayashi. GTRECS2: A model checker for recursion schemes based on games and types. A tool available at <http://www-kb.is.s.u-tokyo.ac.jp/~koba/gtrecs2/>, 2012.
- [21] R. P. Neatherway, S. J. Ramsay, and C.-H. L. Ong. A traversal-based algorithm for higher-order model checking. In *ICFP*, pages 353–364, 2012.
- [22] C.-H. L. Ong. On model-checking trees generated by higher-order recursion schemes. In *LICS*, pages 81–90, 2006.
- [23] C.-H. L. Ong and S. J. Ramsay. Verifying higher-order functional programs with pattern-matching algebraic data types. In *POPL*, pages 587–598, 2011.
- [24] Luke Ong. Higher-order model checking: An overview. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 1–15, 2015.
- [25] S. J. Ramsay, R. P. Neatherway, and C.-H. L. Ong. A type-directed abstraction refinement approach to higher-order model checking. In *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*, pages 61–72, 2014.
- [26] S. Salvati and I. Walukiewicz. Simply typed fixpoint calculus and collapsible pushdown automata. *Mathematical Structures in Computer Science*. To appear.
- [27] Sylvain Salvati and Igor Walukiewicz. Recursive schemes, krivine machines, and collapsible pushdown automata. In *Reachability Problems - 6th International Workshop, RP 2012, Bordeaux, France, September 17-19, 2012. Proceedings*, pages 6–20, 2012.
- [28] S. Schwoon. *Model-checking Pushdown Systems*. PhD thesis, Technical University of Munich, 2002.
- [29] Micha Sharir and Amir Pnueli. *Two approaches to interprocedural data flow analysis*, chapter 7, pages 189–234. Prentice-Hall, 1981.
- [30] D. Suwimonterabuth, S. Schwoon, and J. Esparza. Efficient algorithms for alternating pushdown systems with an application to the computation of certificate chains. In *ATVA*, pages 141–153, 2006.
- [31] Hiroshi Unno, Naoshi Tabuchi, and Naoki Kobayashi. Verification of tree-processing programs via higher-order model checking. In *APLAS*, 2010.